SPECTRUM SYNTHESISER

★ Low Cost
★ Easy to Construct
★ Four Octave Keyboard
★ Performance Controller
★ FM and Sync.
★ Sequencer Effects
★ Stereo Outputs
★ Interface Facilities

by Chris Jordan

The Spectrum is a monophonic two oscillator switch-linked synthesiser featuring advanced specification, constructional simplicity and low cost. Modulation, timbre control, and interface facilities not found on any comparable synthesiser make it extremely powerful and versatile for keyboard playing, sound effects and many other home, stage, or studio applications. Construction is simplified by the use of integrated circuits that each perform major synthesiser functions with few external components. Error-prone control wiring is eliminated by the use of a single PCB mounted behind the panel and holding the pots (PCB mounting) and switches. No glueing of contact blocks or bending of gold wires is needed to assemble the keyboard contacts — a new contact system only requires soldering of the contacts and drilling of the chassis to mount the contact PCB. This also contributes to the low cost — the Spectrum can be built for around £200 including metalwork and PCB's, but not including the case.

Description

Figure 1 shows a block diagram of the synthesiser and the front panel legending is reproduced below. Modulation routing is accomplished by source and function switches and depth controls, rather than the usual method of providing each source with its own depth for each controlled function found on some small synthesisers. Switching is most suitable for a large number of sources as here, and allows fast selection of source and selection of modulation effects with preset depths, in favour of simultaneous modulation of one parameter by more than two signals. Six modulation signals are available: keyboard, controller, low frequency oscillator (LFO), noise generator and external. The keyboard is of the highest note priority type and has a glide which always completes even after the key is released — this makes the keyboard much more useful as a controller for effects sounds. The joystick controller routes a voltage dependent on the side to-side position of the stick to various voltage controlled circuits, allowing it to be used to control the pitch (pitch bend), timbre, or both (see later). The external voltage fed into the controller jack can override or add to the joystick voltage for control by additional synthesiser equipment, or a pedal can be plugged in and used for control by attenuating a fixed joystick voltage.

The low frequency oscillator generates random and regular sample and hold effects in addition to the four common waveforms. The regulator S/H option allows rising and falling scales, rising and falling repeating groups of two, three or more notes, and other sequencer-like effects, with the pattern controlled by the LFO rate. An LED displays the LFO cycle and the joystick's vertical position determines the amplitude at the LFO manual output. The envelope generator is of the exponential ADSR type and like the LFO has + and - outputs that can be separately selected for each controlled parameter. The envelope generator shares its gate signal with the envelope shaper, which determines the loudness contour of each note. 'Single' on the gate selector switch causes gating each time a first key is depressed; 'Multiple' retriggers when any new note is played, allowing fast runs without 'missed' notes. 'Hold' keeps the gate high for continuous effects, and 'LFO' causes gating on each LFO cycle. In the 'Repeat' position the envelope generator retriggers at the end of the decay period, acting as an additional LFO with variable symmetry. This allows complex rhythmic

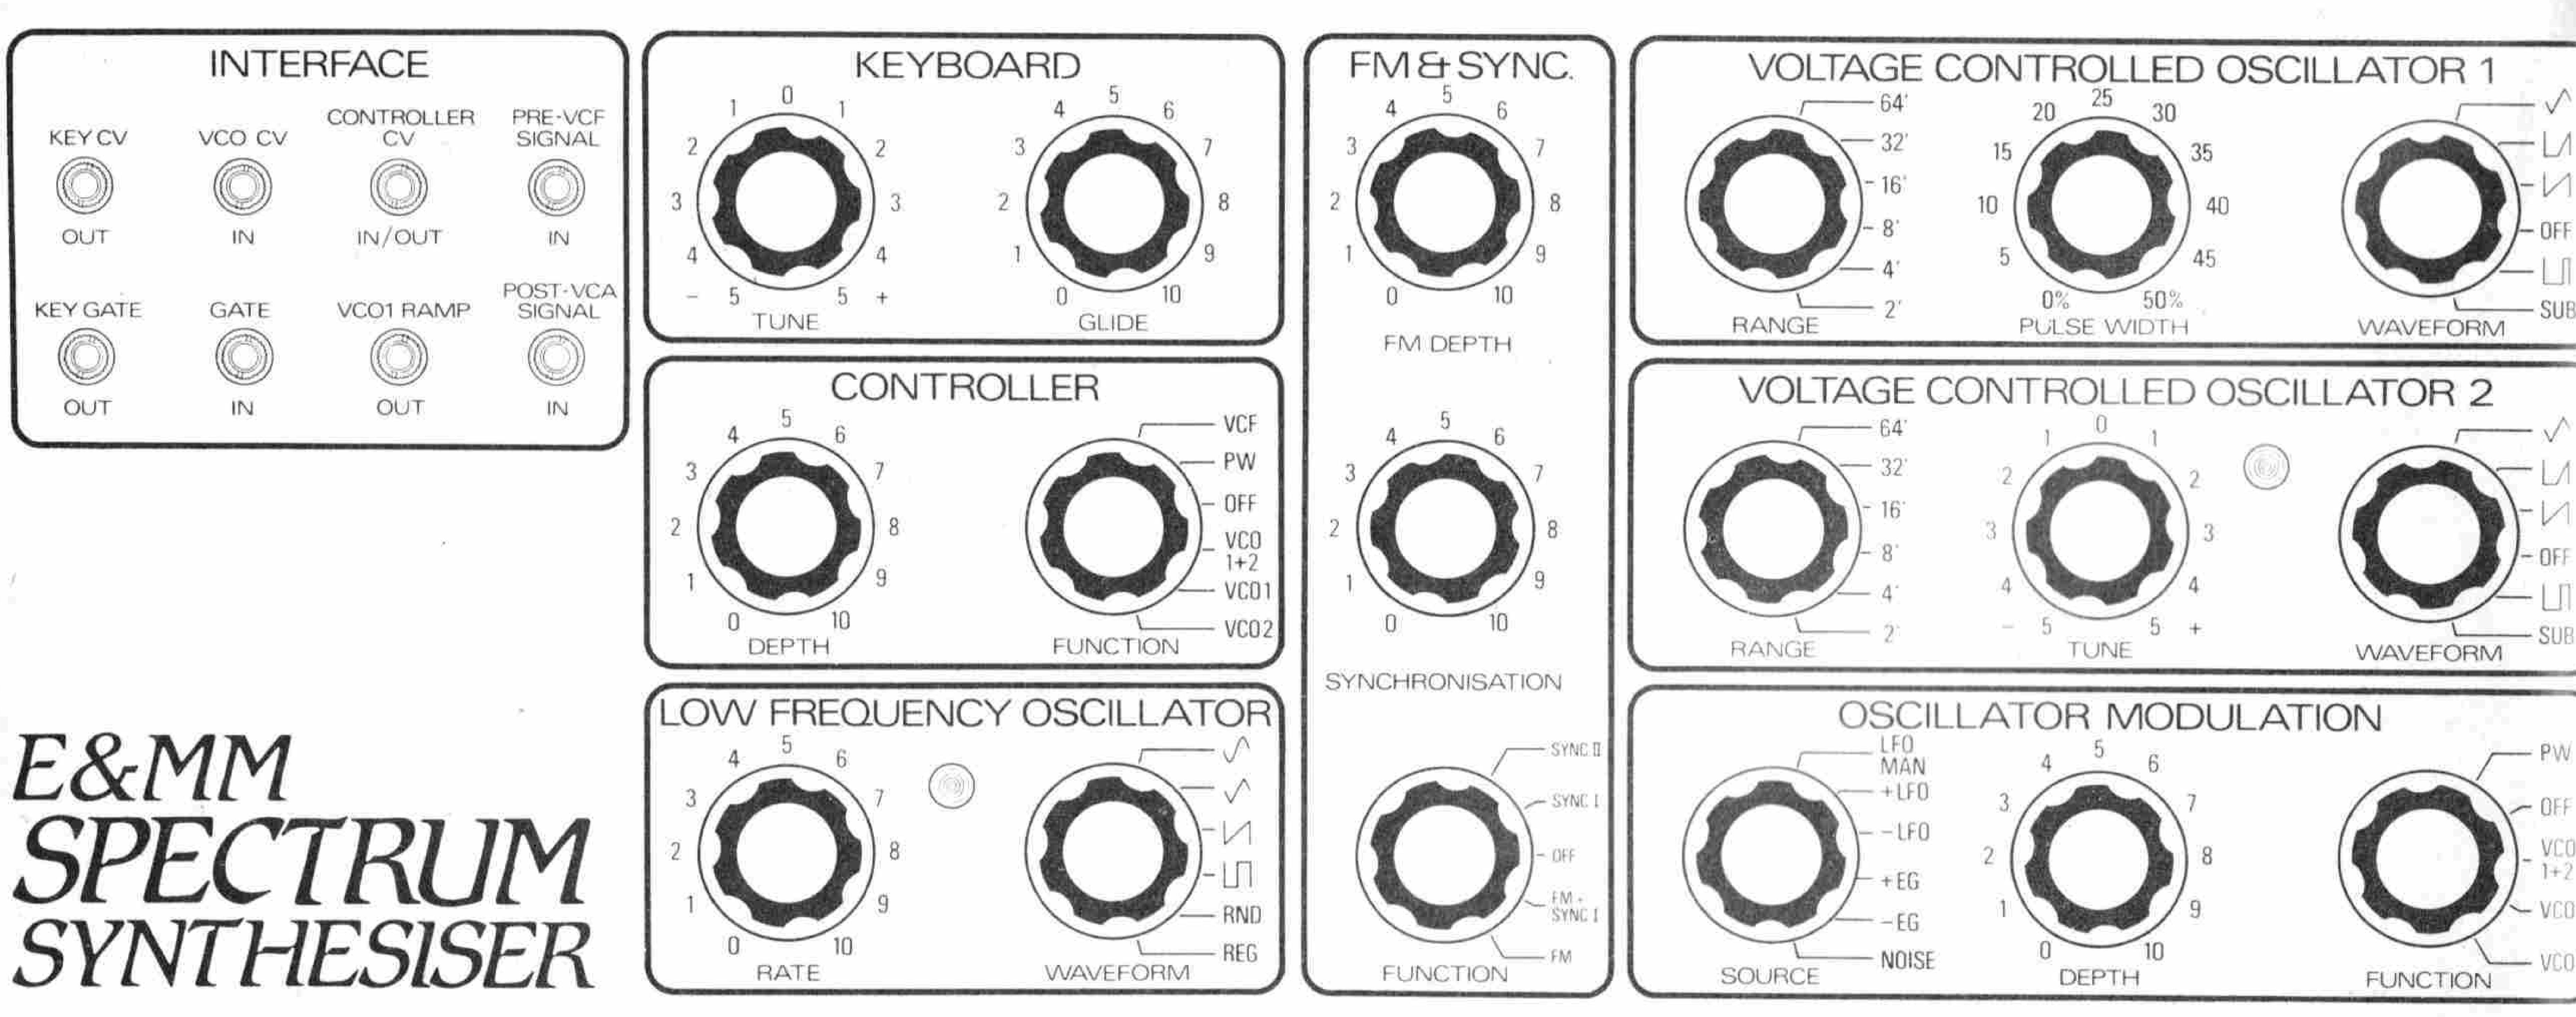

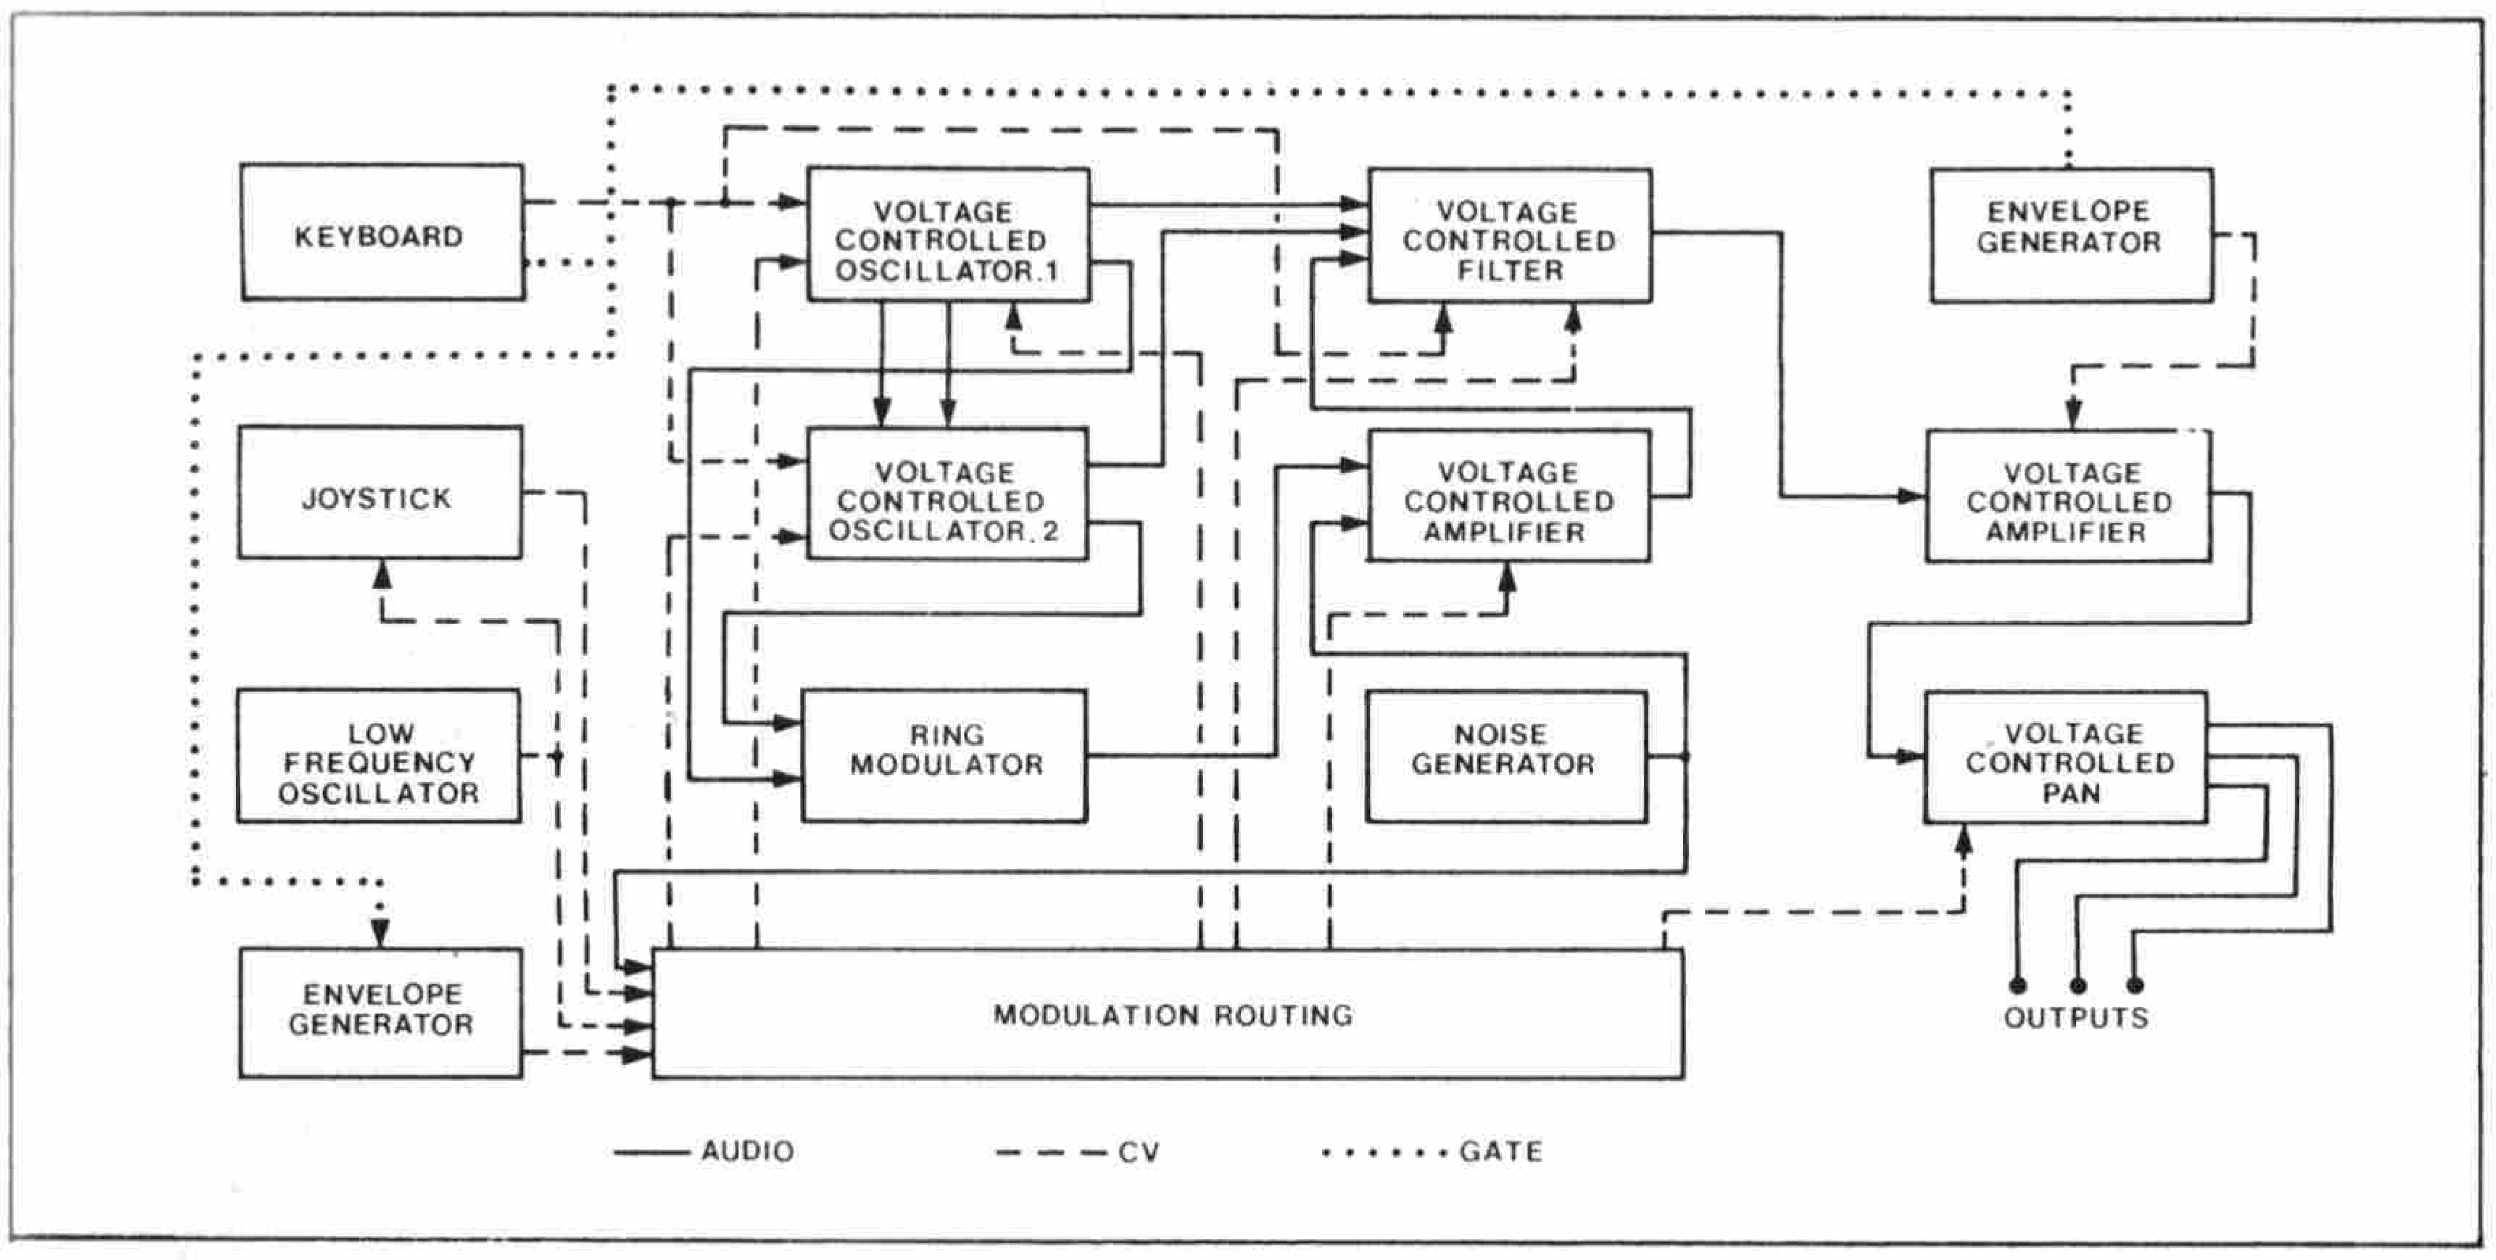

Figure 1. Block Diagram of the Spectrum Synthesiser.

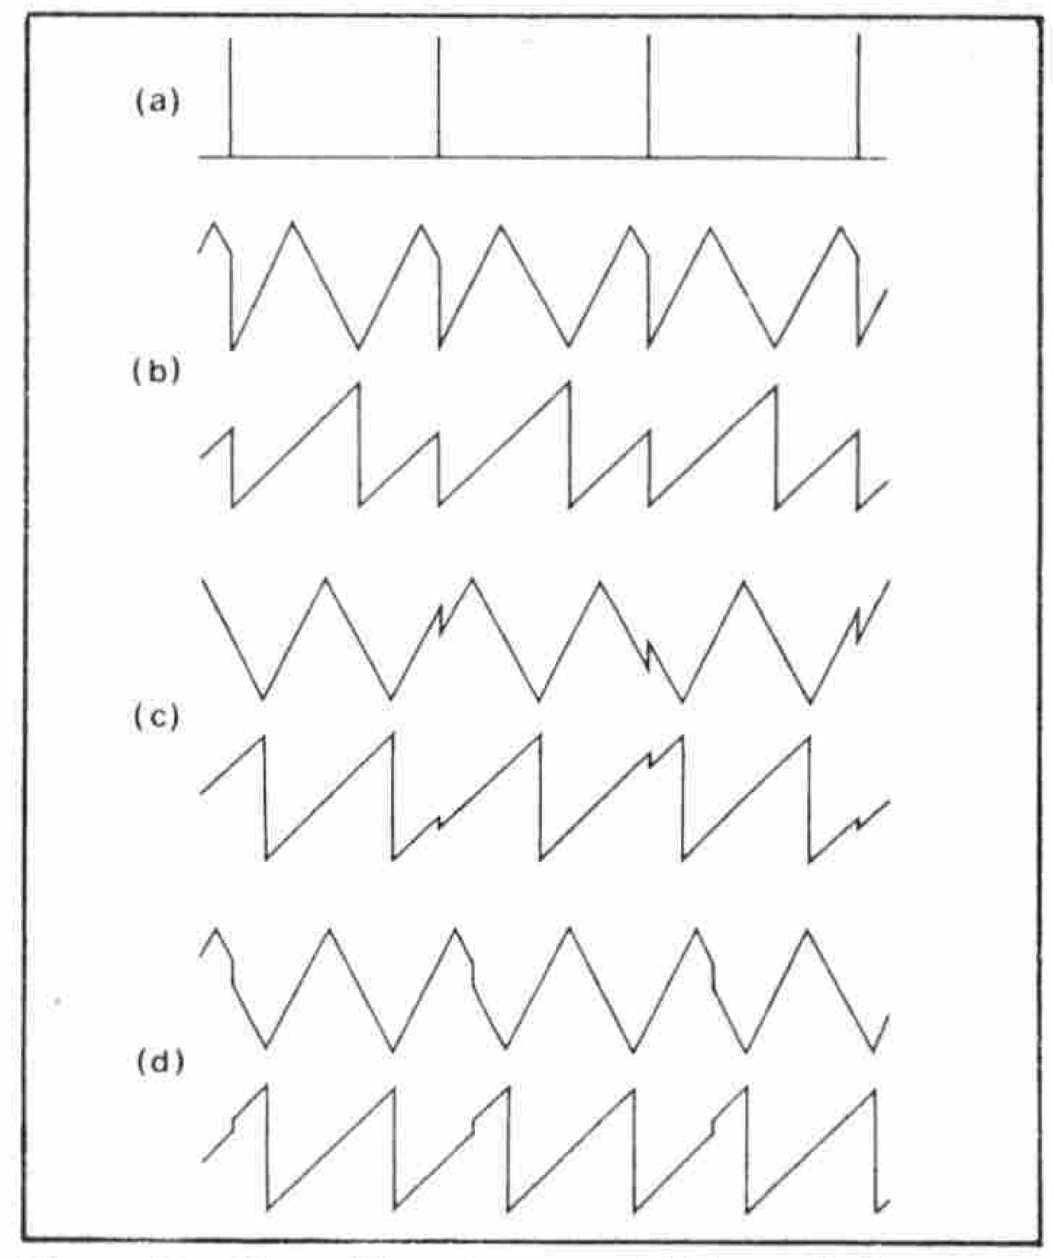

Figure 2. Sync. Waveforms. (a) Sync. Pulses. (b) Sync. I. (c) Sync. II. (d) Sync. II with decreased VCO1 frequency.

effects when used with the LFO, and gives great scope for 'backdrop' sounds based around complex S/H patterns with periodic timbre sweeping effects derived from the EG. 'Key Repeat' brings in the repeat only when a key is held, allowing key-synchronised repeating notes and delayed modulation (the delay determined by the attack time). An LED indicates the EG's attack segment.

The voltage controlled oscillators (VCO's) each have five switched octave ranges and five waveforms. The sub-octave output is a pulse wave with a square wave added an octave below, making the sound fuller and richer. The tuning LED detects the beats between the oscillators, and indicates when the pitches are in simple musical intervals, useful for tuning without sounding a note (e.g. on stage). The pulse width of VCO 1 is variable, and VCO 2 has a tune control with a ± one fifth range.

The VCO's can be used together to provide a vast range of sounds not possible with basic synthesisers having only waveform, shape, VCF cutoff and VCF resonance as the controls affecting basic timbre. This is done by frequency modulation and synchronisation — special features of this design. FM uses the triangle output of VCO 1 to modulate the frequency of VCO 2 up to ±100% giving a whole range of non-harmonic tones for bell, gong, chime sounds etc. Synchronisation gives various waveforms from VCO 2 (see Figure 2) which have particular bands of harmonics emphasised for strong, voice-box-like sounds. This is achieved by resetting the output of VCO 2 upon each cycle of VCO 1, so the tones generated are always harmonic. Two modes of sync. are provided: Sync. I is that normally found on rampwave oscillators, the VCO 2 waveform beginning in the same way after each reset; Sync. II is something totally new — the triangle output is set to mid way each time but then carries on in the same direction in the new cycle. VCO 2 locks on to VCO 1 harmonics with the change from one harmonic to the next emphasised by a sharp change in tone. This enables automatic arpeggiation and incredible tone sweeps to be obtained since VCO 2 now is effectively a voltage controlled waveform generator/frequency multiplier. The sync. control attenuates the pulses fed to VCO 2 so that it only resets if the wave form is above a certain threshold, resulting in the oscillators being locked together in musical intervals (3rds, 5ths etc.). Simultaneous sync. I and FM produces harmonic tones with the shape of FM-ed waveforms within each cycle.

The ring modulator uses triangle and square VCO waveforms to provide further complex tones. Its output is mixed with the noise signal and fed into a special voltage controlled amplifier (VCA). This can be controlled by the LFO or EG, and gives the signals their own loudness contours. Hence noise 'chiffs' can be added to notes, or ring modulation set to swell in as a note decays.

The VCA output is fed to the voltage controlled filter (VCF) mixed with the VCO outputs. The VCF offers the two most useful responses, low pass and band pass, plus an intermediate response for bright sounds that remain strong in lower harmonics. Cutoff frequency and resonance controls perform their normal functions and a keyboard follow control determines how the cutoff frequency varies over the keyboard range.

After envelope shaping, the signal is fed to the voltage controlled pan circuit which can modulate the location of the sound in the stereo field by the LFO or EG signals. The stereo outputs can also be used for

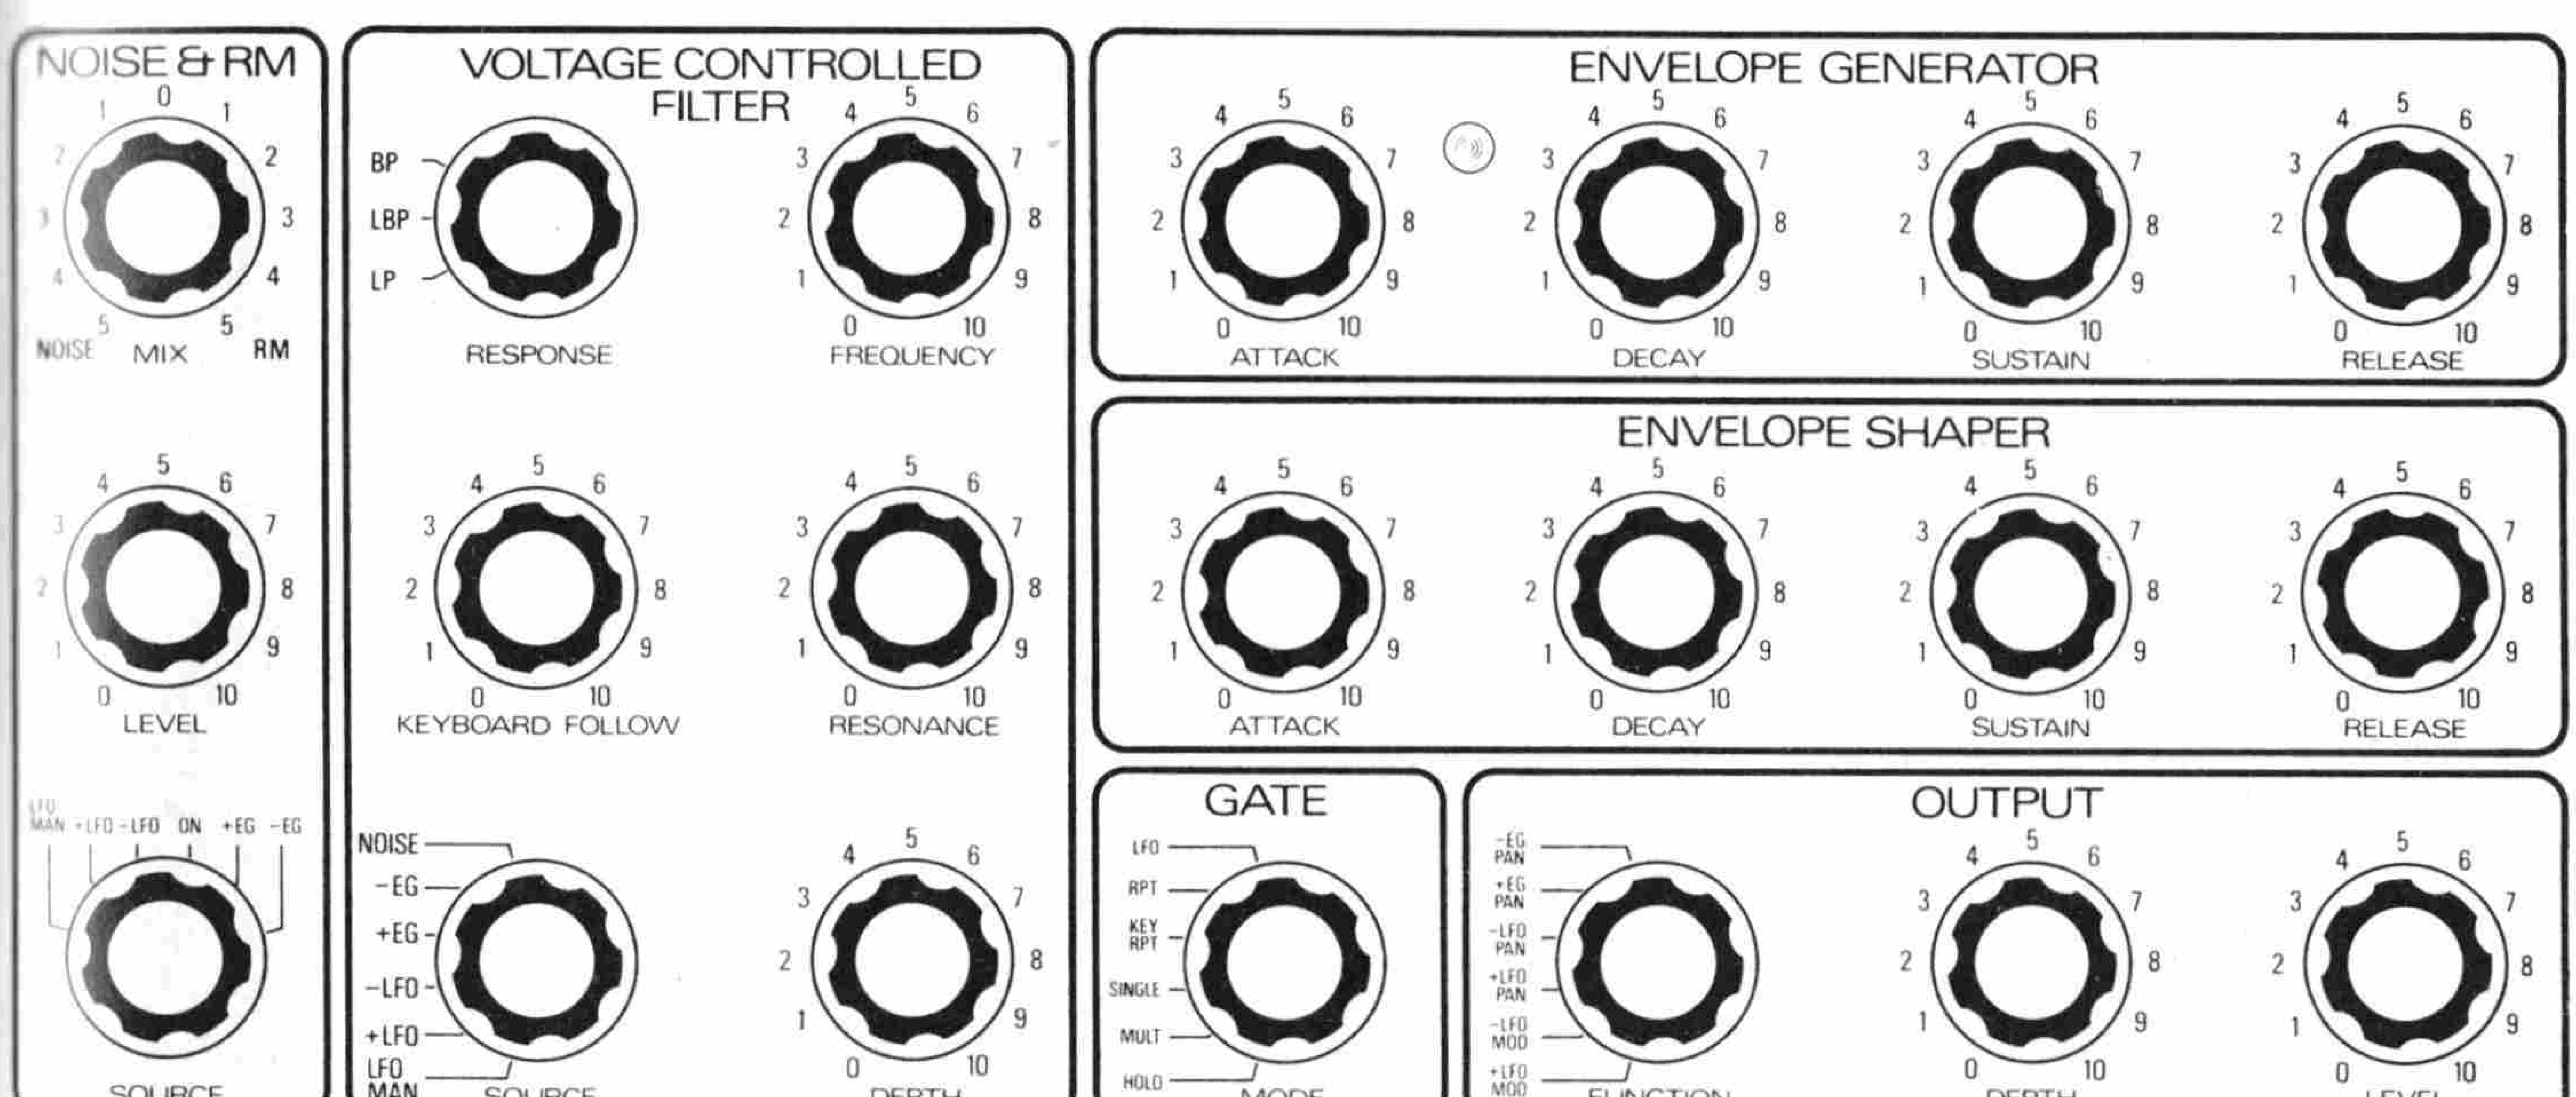

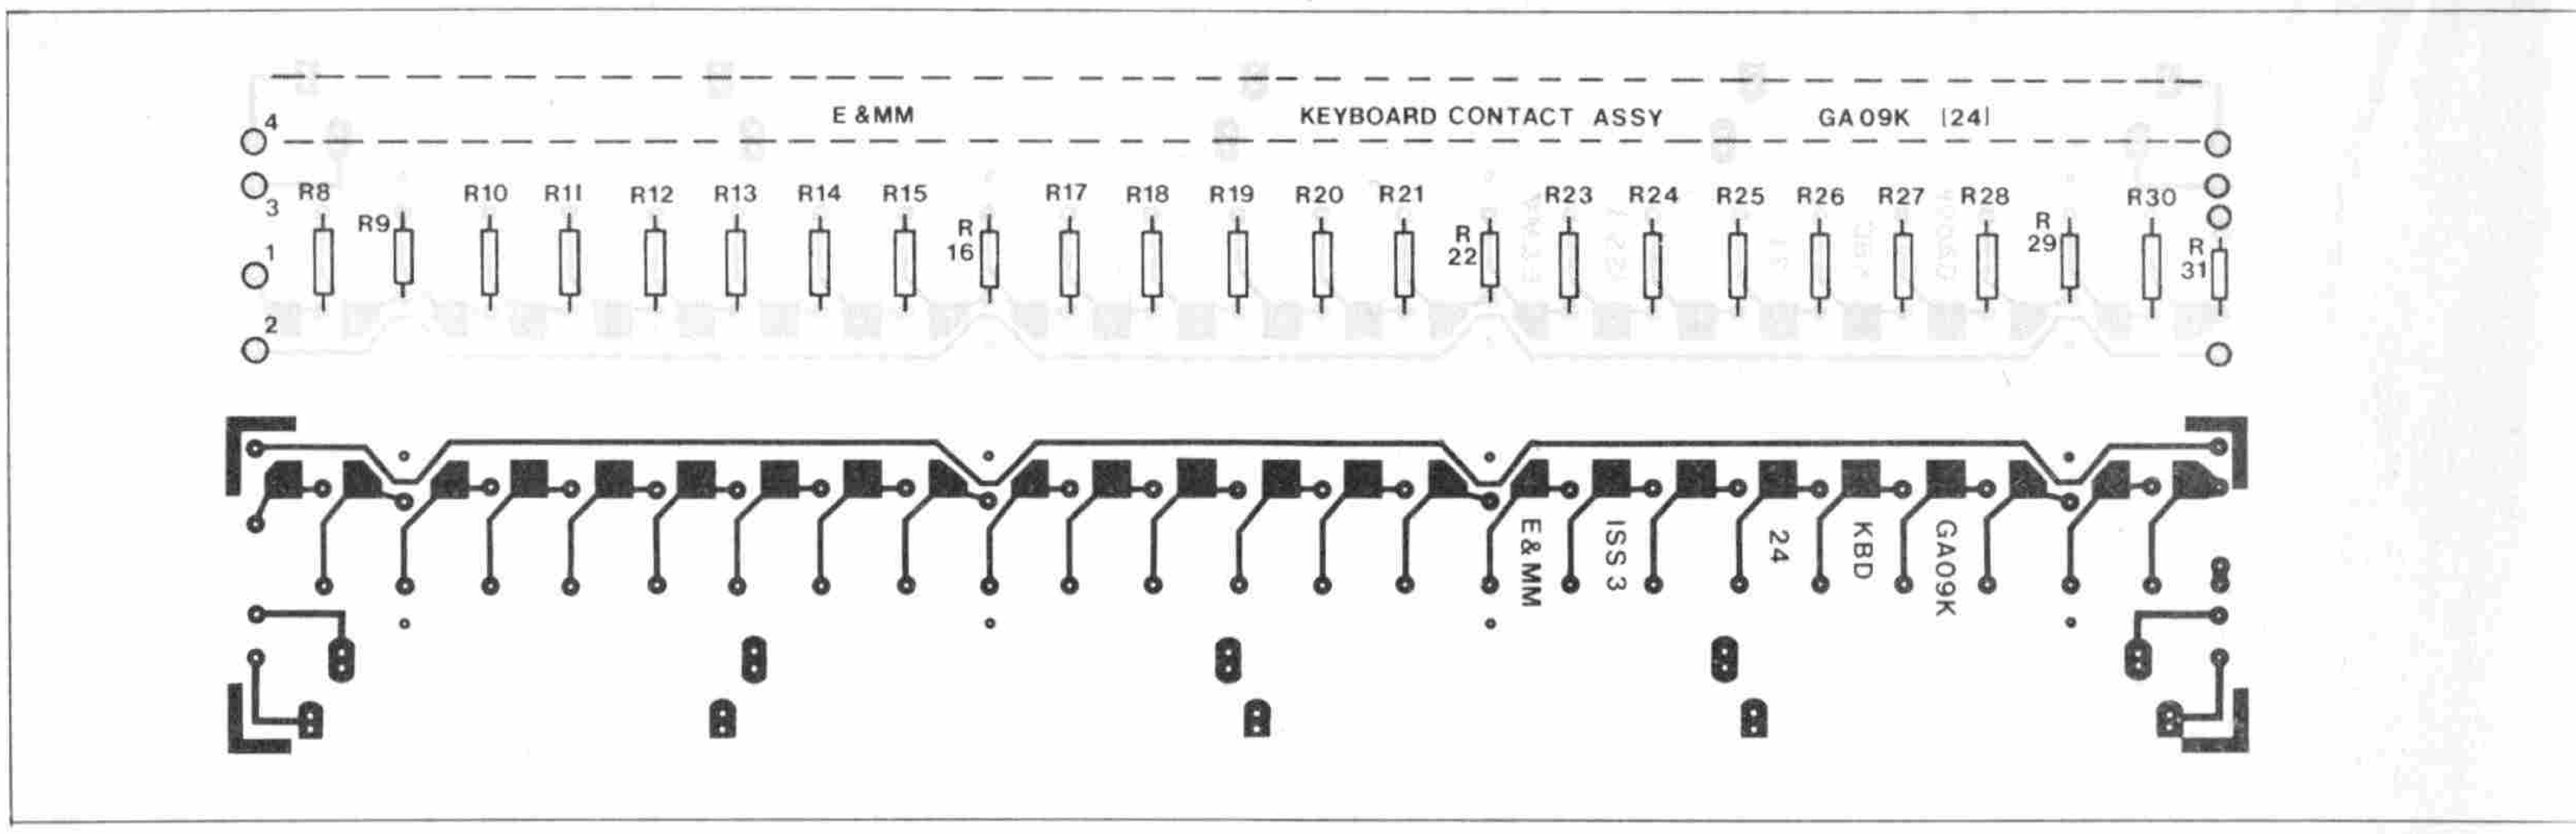

Figure 5. Keyboard printed circuit board.

voltage control of the depth of external effects such as reverb, phase, and echo, by routing one signal via the effects unit and one direct to the amplifier. A mono output is also provided, and the pan VCA can also be used for additional amplitude modulation with the LFO as source (for tremolo and other effects).

The interface jacks allow connection to external devices such as sequencers, additional VCO banks, waveform processors etc, and designs for these will be published in the future. The Spectrum Synthesiser and future equipment will use the 1V/octave CV standard, and can be interfaced to any other exponential CV synthesiser.

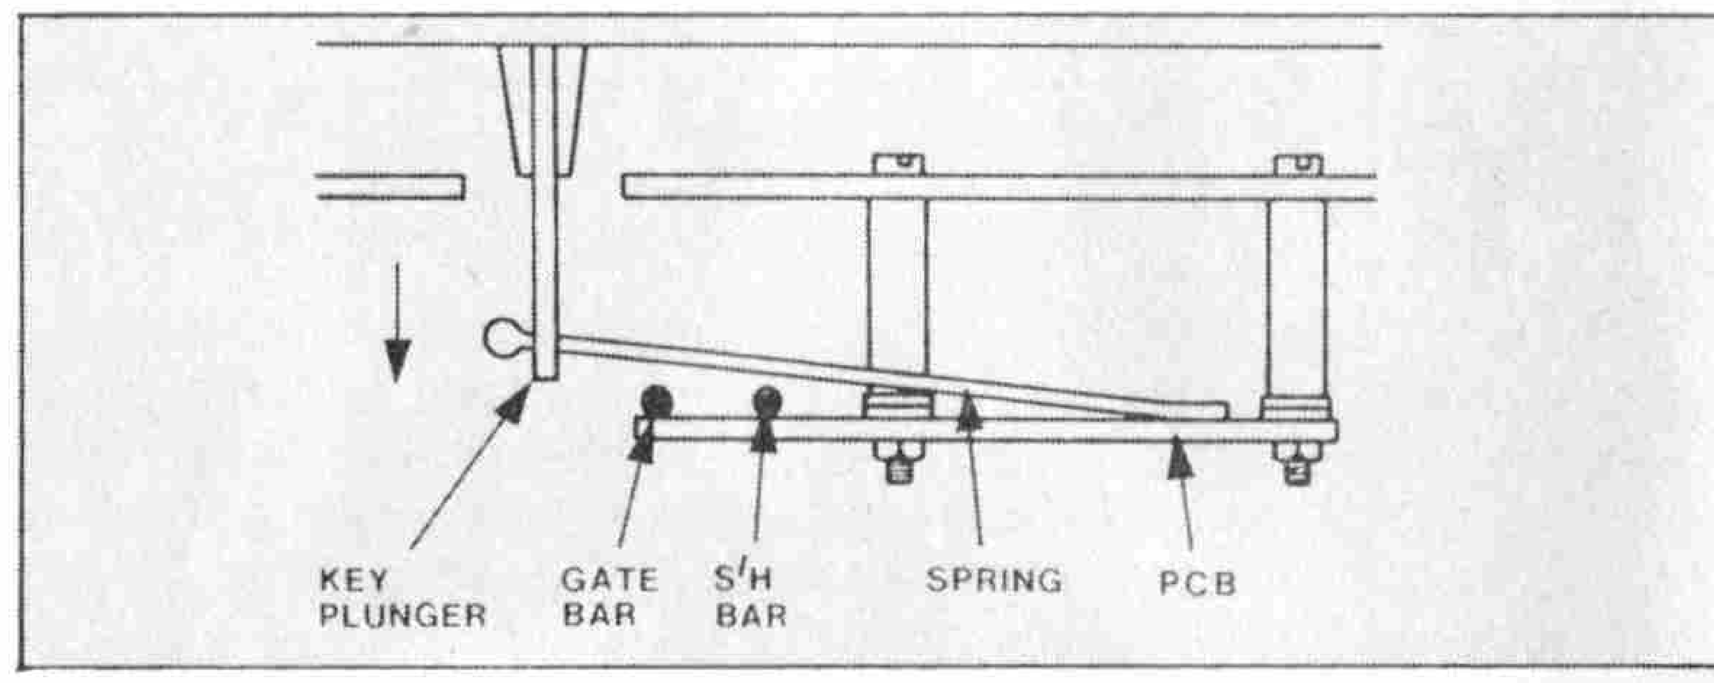

Figure 3. Key contact construction.

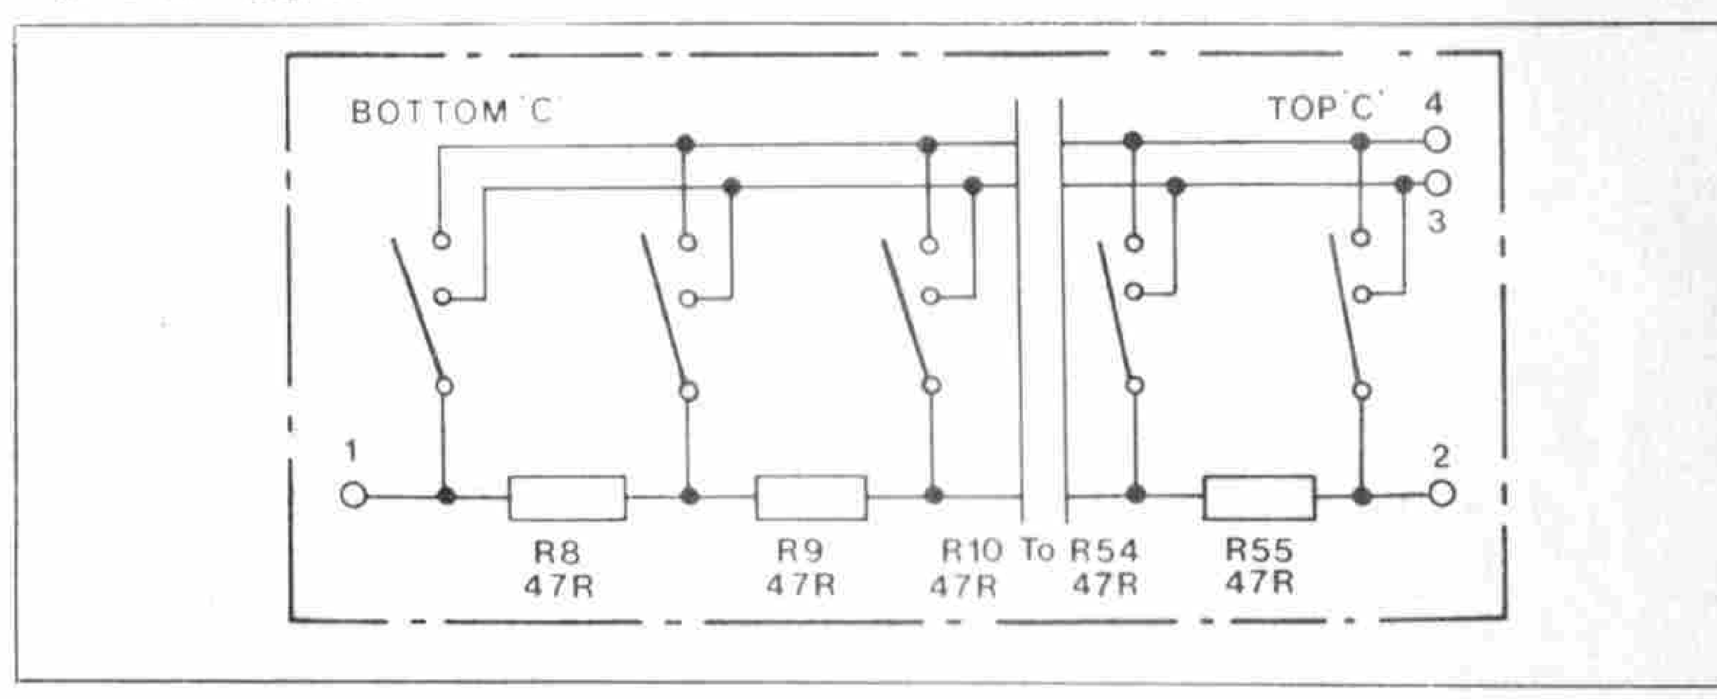

Figure 4. Circuit of key contact assembly.

Keyboard

The Spectrum uses a unique key contact system which is cheaper, more reliable, and easier to construct than alternatives using gold-plated wire and contact blocks. A single moving contact is used for each key with all contacts and their associated divider chain resistors held on a PCB (in two parts) fixed to the keyboard chassis.

The moving contacts are silver plated springs, each fixed at one end and moved at the other by the plunger of the respective key such that the spring makes contact with two palladium bars when the key is depressed (Figure 1). The first bar is connected to the sample and hold circuit which stores the voltage representing the last key depressed, and the second to a circuit which generates a gate signal for the S/H and the envelope generators. The moving contacts connect to the divider chain (see Figure 2). These functions are usually carried out by separate contact pairs, where unless the contacts are precisely set up, note-jumping will occur when the envelope is gated before the S/H receives the new key voltage. The system used here is immune from this since the construction ensures the correct sequence of operation, and no initial setting up is required. The keyboard recommended in the parts list has removable key plungers so that cleaning the contacts is much easier too. Unclipping a plunger allows access to the sides of the bars and springs that meet.

Keyboard construction

Use the printed circuit board as a template to mark the fixing holes on the underside of the keyboard chassis. Mark them such that the edge of the board holding the bars will be about 5mm from the plungers and then drill for 6BA clearance. Fit the 48 divider resistors on the component side of the board along with the 12 veropins and solder in place. Cut the palladium bars to length and fit them to the track side using small loops of wire passed over the bar, through the mounting holes and twisted on the component side. Make sure each bar is well seated before soldering at each loop position on both sides. Cut each plunger to length, leaving the nearest slot to the key end for the contact. Tin 5mm of both ends of the contact springs and fit each one by passing the thin end through the detached plunger and soldering it to the pad on the PCB. If you've marked the PCB mounting holes correctly then for proper operation the end of the spring should be about 2mm from the far edge of the pad. The positioning of the PCB and the springs on the PCB is not critical as long as when the PCB is mounted and the plungers clipped on, the springs are under slight tension to ensure positive contact. Mount the PCB to the chassis using 6BA bolts, ½" spacers and nuts, and washers to separate them further. The keys opposite the mounting positions will have to be temporarily removed to fit the bolts, and this should be done before drilling if a hand-held drill is used, to avoid the possibility of damage to the keys. Again, the spacing is not critical so long as all the contacts normally clear both bars and make contact with both when their keys are depressed.

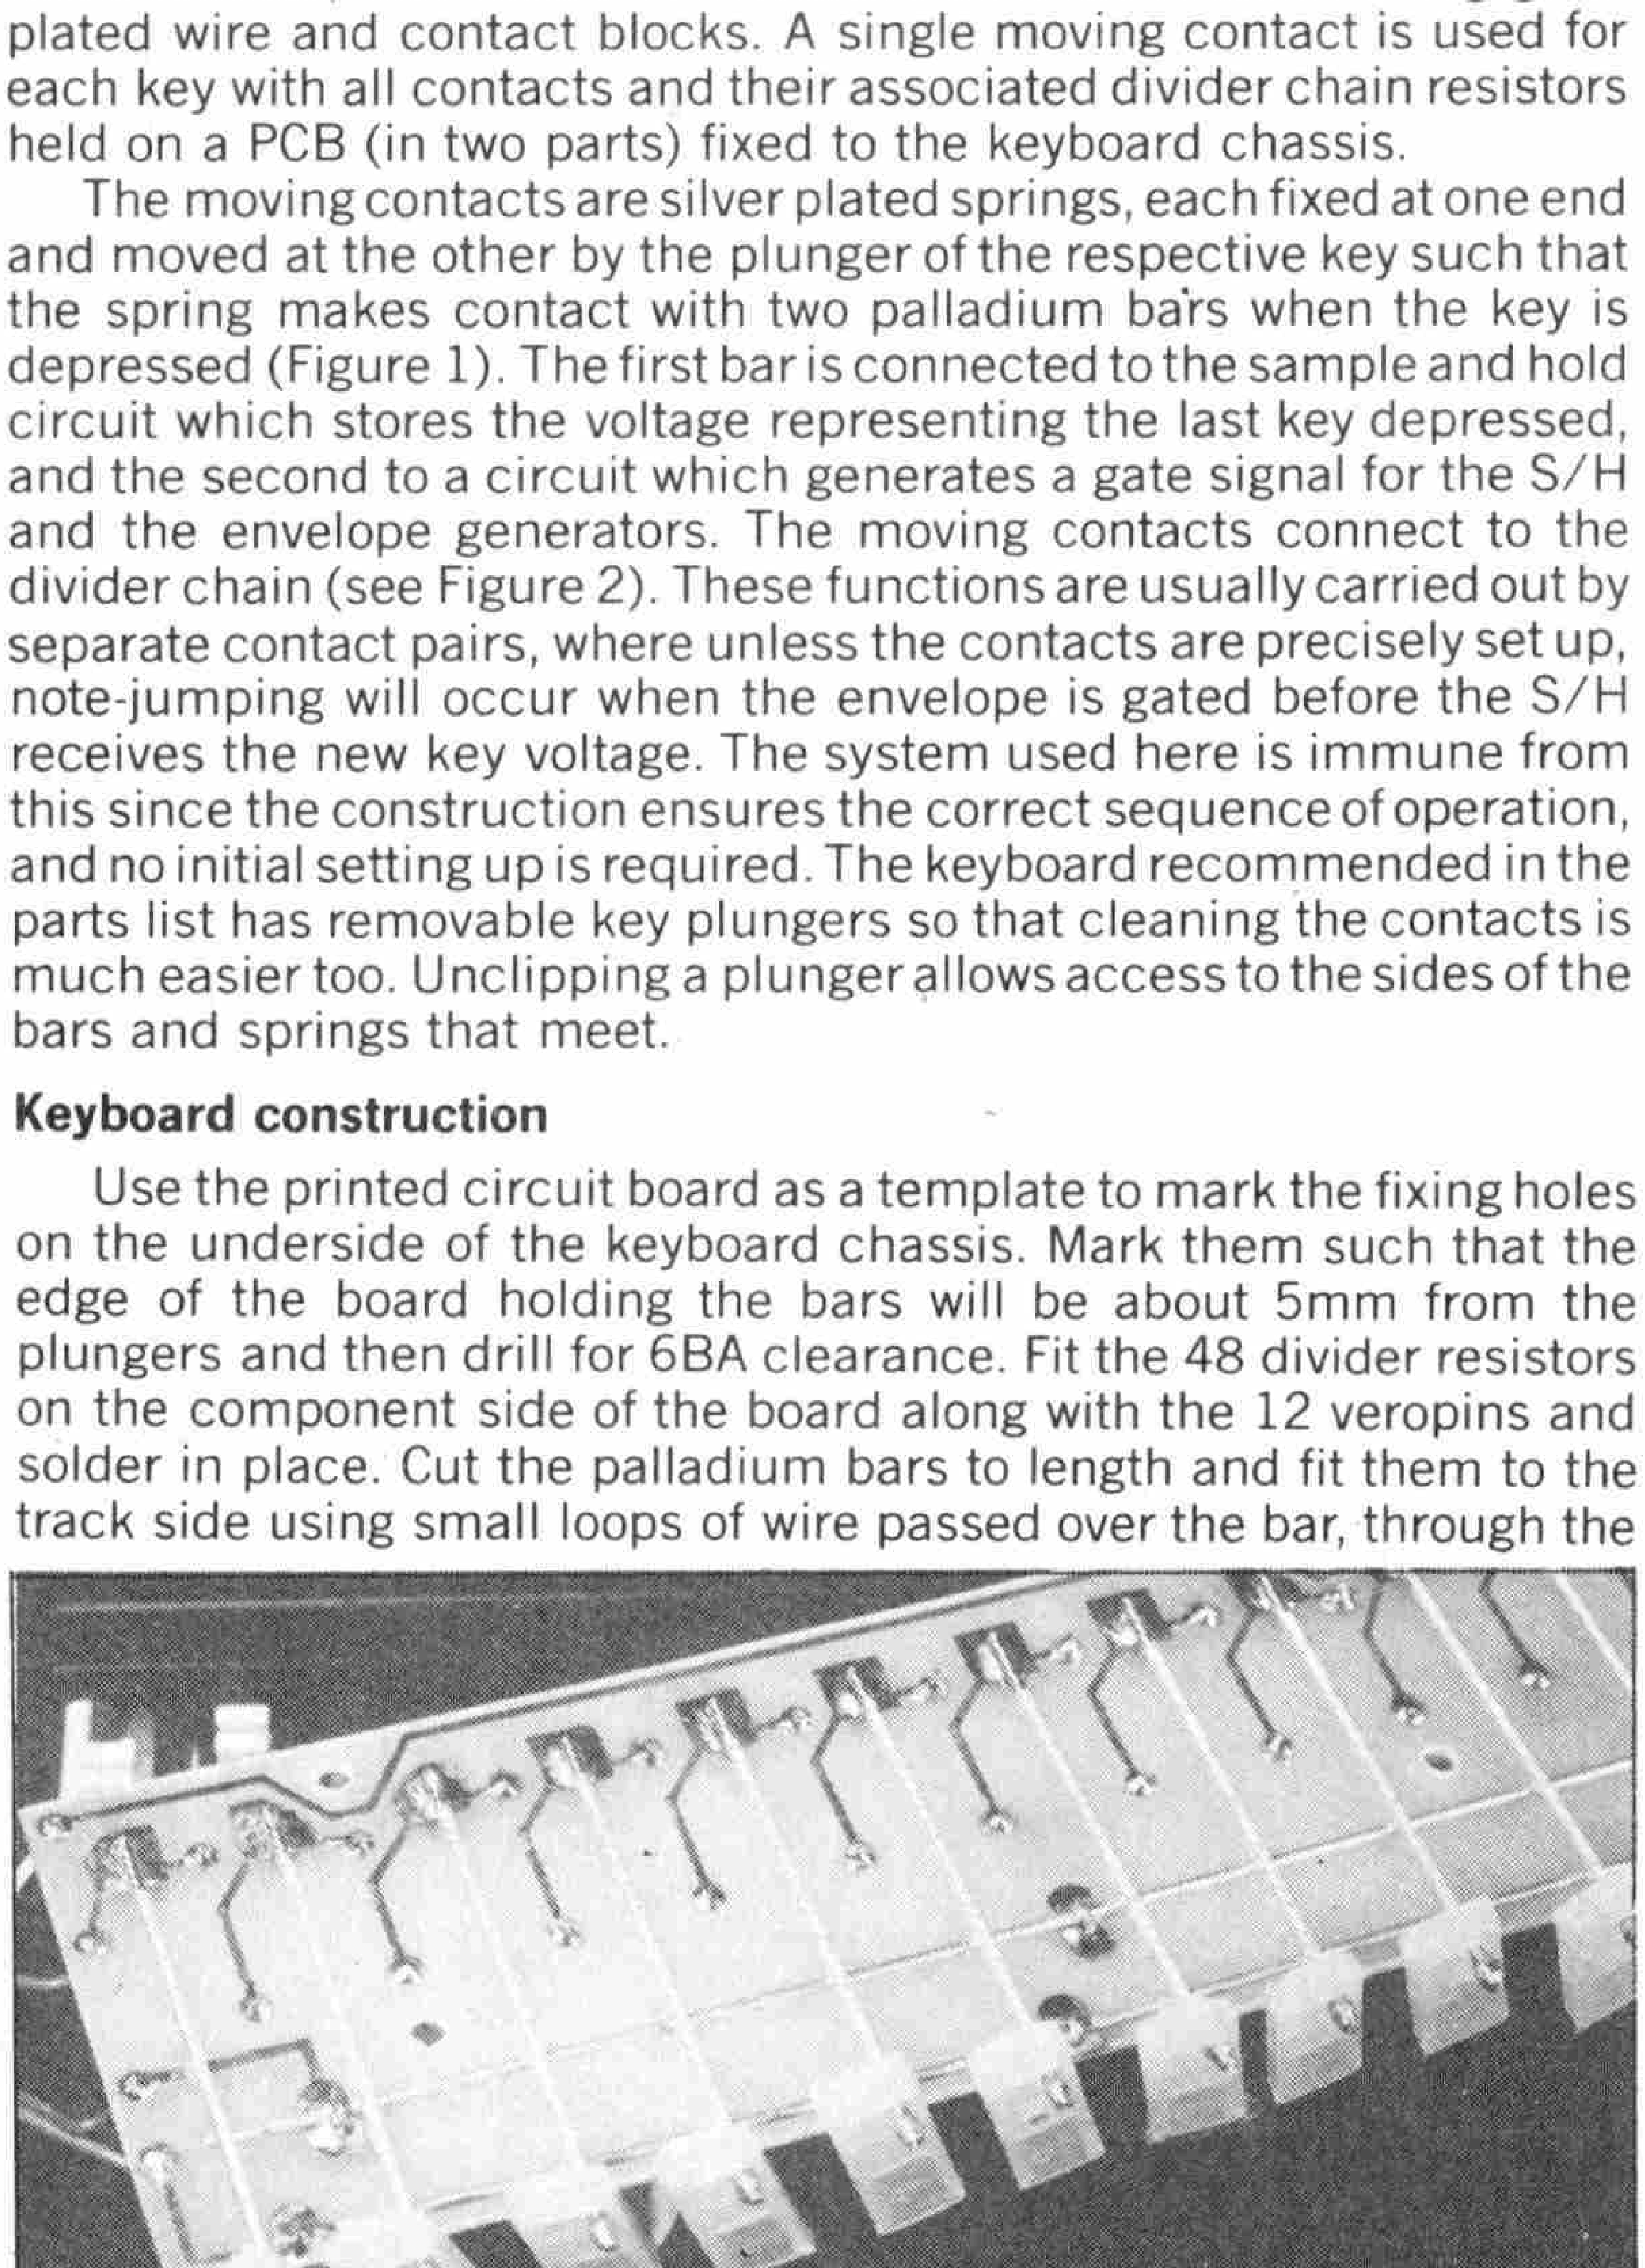

A view of the key contacts before mounting of the board.

KEYBOARD PARTS LIST

Resistors			
R8-55	47R 2%	48 off	(X47R)
Miscellaneous			
	49-note C-C keyboard		(XB17T)
	Contact springs	49 off	(QY07H)
	Palladium bars, 1.2mm × 330mm	Set of 4	
	24-contact PCB		(GA09K)
	25-contact PCB		(GA10L)
	6BA 1" bolts		(BF67H)
	6BA ½" spacers		(FW35Q)
	6BA washers		(BF22Y)
	6BA nuts		(BF18U)
	Veropins		(FL24B)

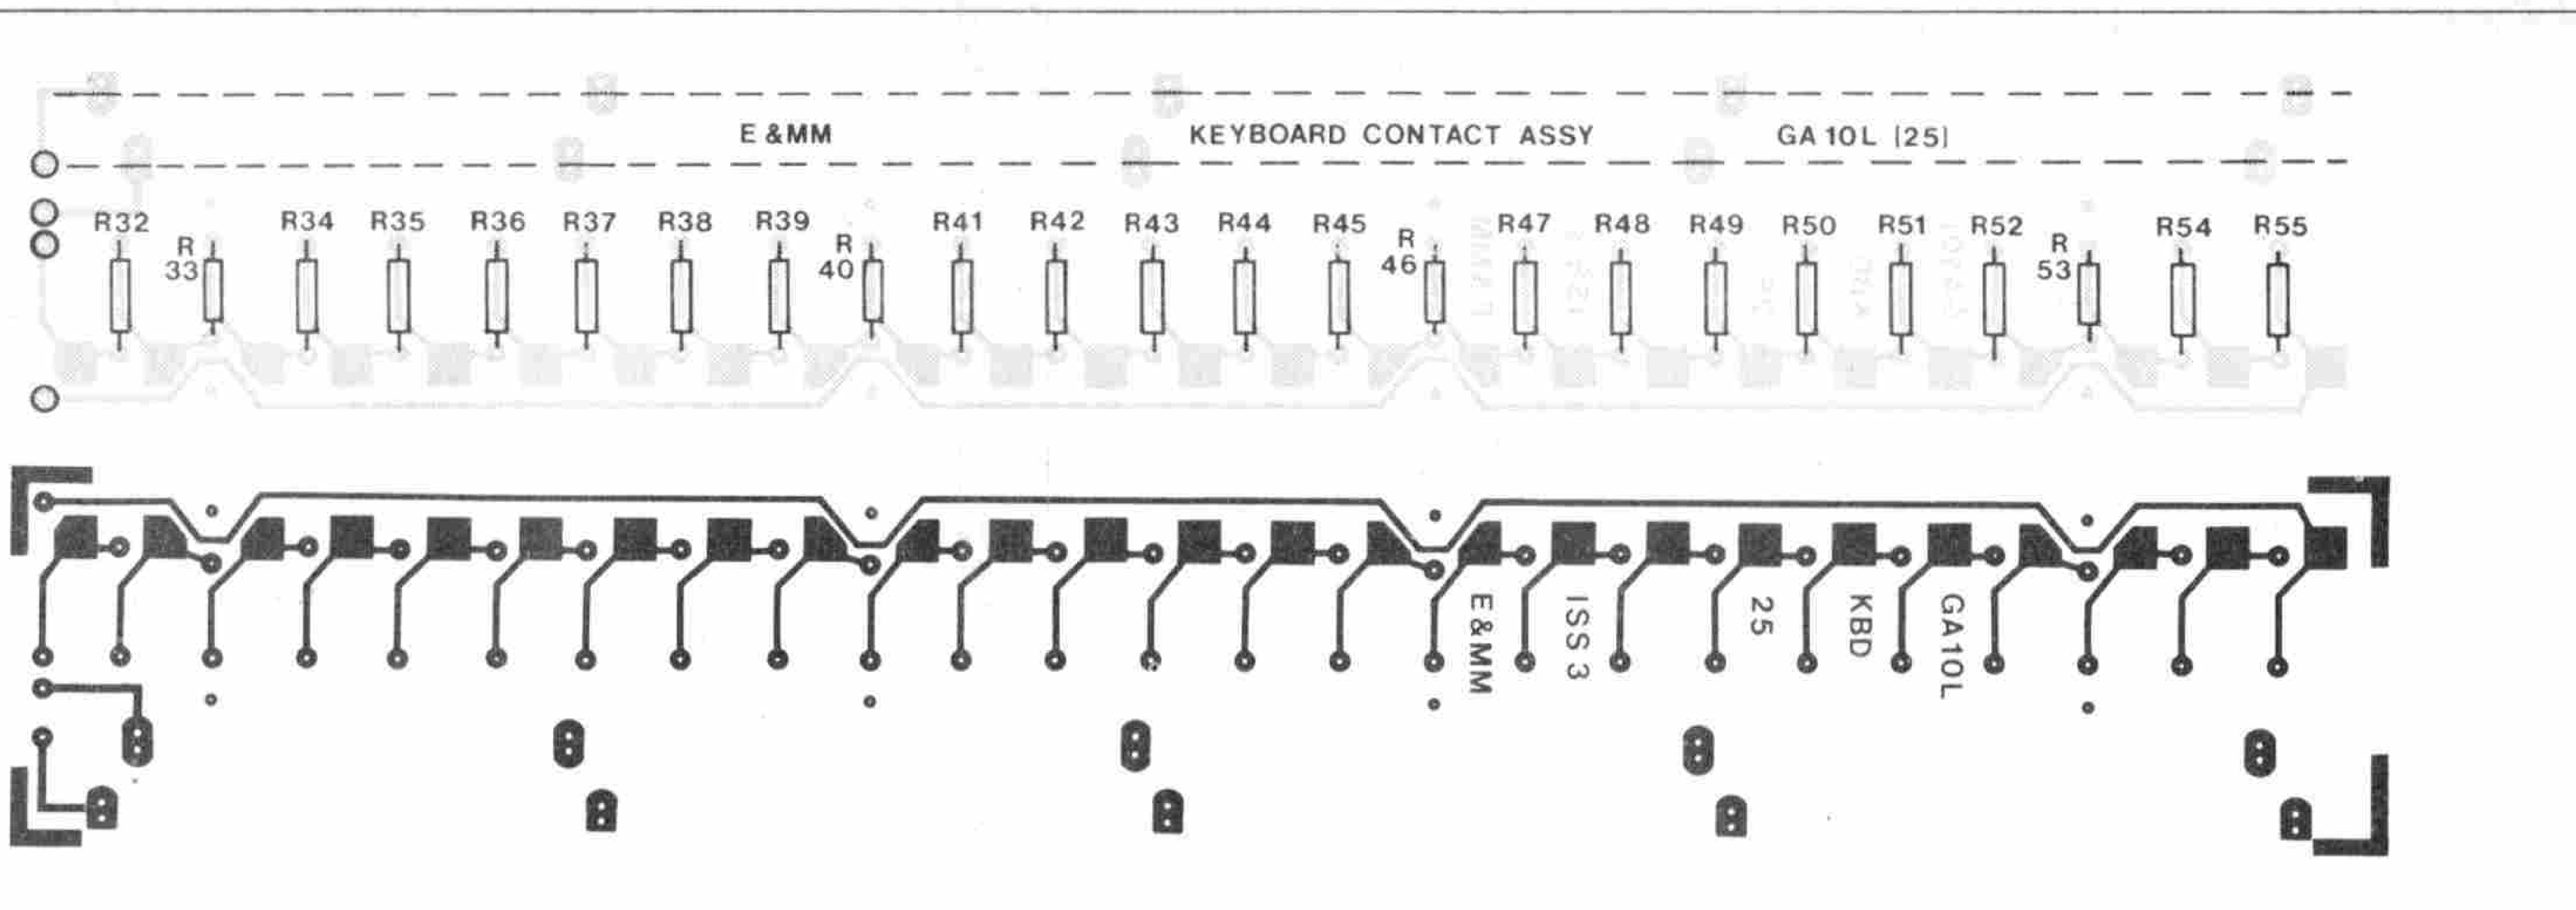

A ½" spacer and one nut were found to be about right, though washers could be used if a high or low action to the keys is preferred. Connect the two halves of the board together using short wire links across the veropin pairs. This completes the keyboard construction.

The assembly can be tested on its own by connecting a multimeter set on the low resistance range across pins 1 and 3. Depressing the bottom key should give zero resistance, and the top key about 2.4K. Check that all the other keys give intermediate readings and repeat for the other bar with the meter across pins 1 and 4.

Power Supply Unit

The proper operation of synthesiser circuits requires a stable, noise-free supply, so it is important that the Power Supply Unit (PSU) is well-regulated and has current in reserve. The Spectrum PSU uses monolithic regulators and low temperature coefficient components in a dual design to provide ± 15V at 270 mA maximum.

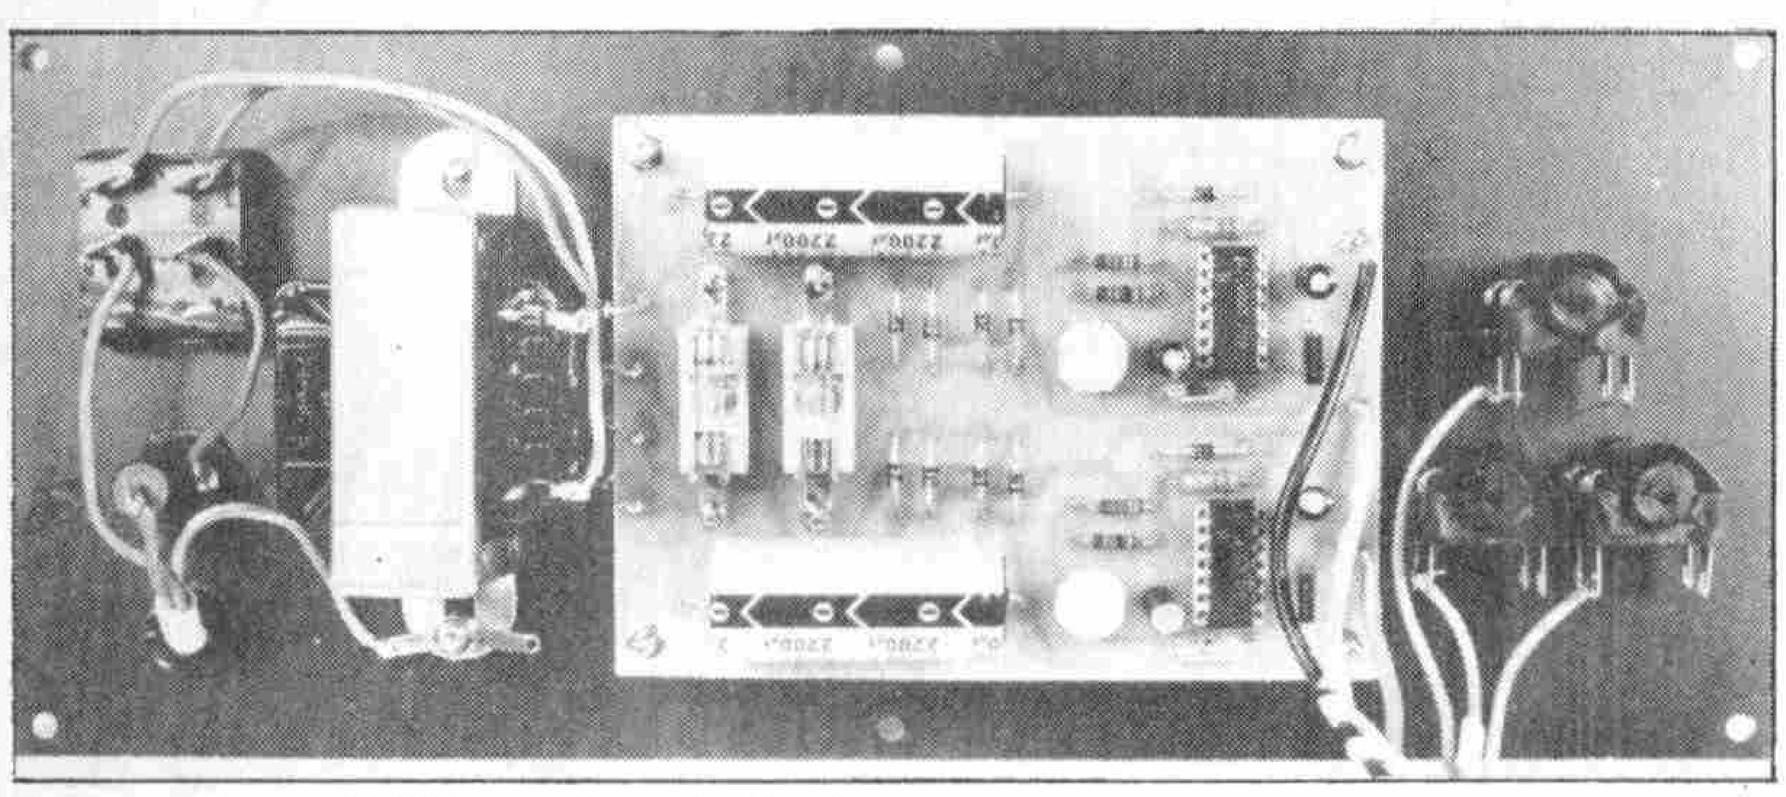

The completed PSU fitted to the back plate.

PSU Circuit

The Power Supply Unit consists of two identical circuits providing the positive and negative supplies, driven by a dual secondary transformer. Each secondary produces about 21V when the AC signal is rectified and smoothed, and is fused for protection in the event of a power supply fault. Regulation is carried out by the well-known uA723 regulator IC which is used with an external power transistor in series pass mode to provide the required current. This current limits at 270 mA when the voltage across series resistor R1 (R2 in the -ve side) reaches 0.6V. RV1 (RV2) allows the rail voltage to be adjusted to exactly 15V, and D1 (D2) protects against reverse polarity, again in the event of a fault. The +15V regulated output of the side based around IC2 is connected to 0V of the IC1 side, giving the -15, 0, +15V supply rails.

PSU Construction

The prototype PSU was mounted on an aluminium panel along with the mains cable, switch, and fuse. This was then fixed by screws to the wooden back of the synthesiser allowing the PSU to be separately assembled and easily removed if necessary (see photograph).

All components except the transformer fit on the PCB. Assemble the components onto the printed circuit board, starting with the resistors, diodes, polystyrene capacitors and IC sockets. These can all be inserted and soldered in together, before the large components are fitted. Note the orientation of the diodes, particularly in the bridge rectifiers (D1-8).

Use 4BA bolts and nuts to fix the chassis fuse holders, and connect them using veropins through the tag holes. Insert and solder in the cermet presets and electrolytic capacitors, followed by the two power transistors. The leads leave the transistors quite close together and should be bent apart about ⅛" from the package before putting them in place. Check that you have put them in the right way round, with the metal sides facing the nearest board edge. Solder in the eight remaining veropins and check the orientation of the electrolytic capacitors diodes, and transistors. Fix the PSU board to the back panel, or whatever else you are using, using 6BA bolts with spacers. Fit the mains switch, fuseholder, and transformer to the panel, using 4BA bolts for the latter and including two solder tags on one side for the

POWER SUPPLY UNIT PARTS LIST

Resistors — all 5% ⅓W carbon unless specified.			
R1,2	2R2 ½W	2 off	(S2R2)
R3,4	3k3 1%	2 off	(T3K3)
R5,6	3k0 1%	2 off	(T3K0)
R7	330R		(M330R)
RV1,2	1k cermet preset	2 off	(WR40T)
Capacitors			
C1,2	2200uF 25V axial elect.	2 off	(FB90X)
C3,4,7,8	2u2 63V PC elect.	4 off	(FF02C)
C5,6	100pF polystyrene		(BX28F)
Semiconductors			
IC1,2	uA723 14-pin DIL	2 off	(QL21X)
TR1,2	BD135	2 off	(QF06G)
D1-D10	1N4001	10 off	(QL73Q)
Miscellaneous			
T1	240V prim. 0-15, 0-15 sec. 10VA		(LY03D)
S1	DPST rocker switch with neon		(YR70M)
FS1	20mm 500mA quick blow fuse		(WR02C)
	20mm panel fuseholder		(RX96E)
FS2,3	20mm 1A quick blow fuse	2 off	(WR03D)
	20mm chassis fuseholder	2 off	(RX49D)
	14-pin DIL socket	2 off	(BL18U)
	Printed circuit board		(GA03D)
	3A 3-core mains cable 2m		(XR01B)
	13A mains plug		(HL58N)
	6BA 1" bolts		(BF07H)
	6BA ½" spacers		(FW35Q)
	6BA nuts		(BF18U)
	4BA ½" bolts		(BF03D)
	4BA nuts		(BF17T)
	4BA solder tags		(BF28F)
	Cable grommet		(LR48C)
	Veropins		(FL24B)

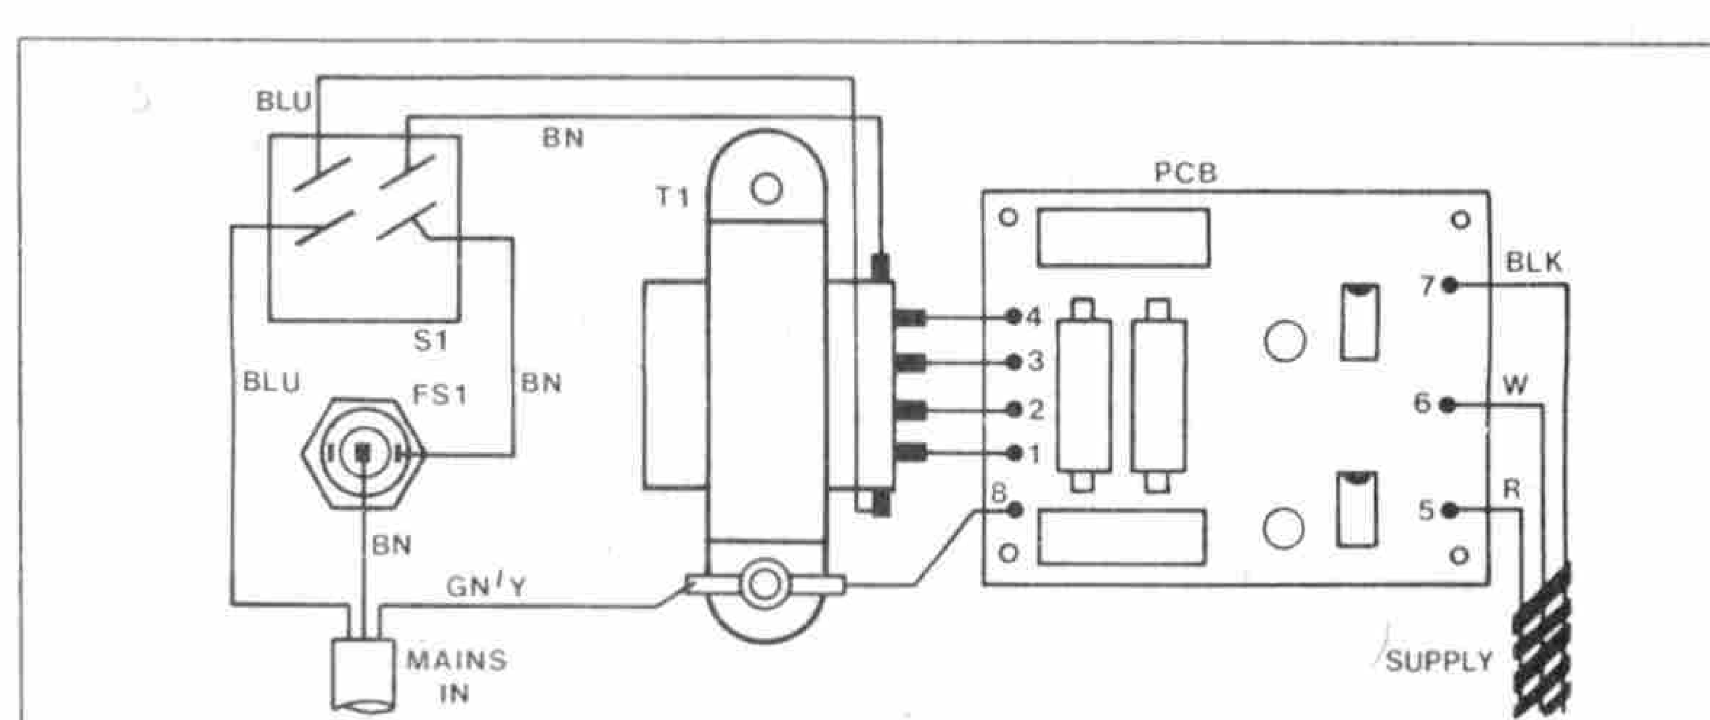

Figure 8. PSU and mains wiring.

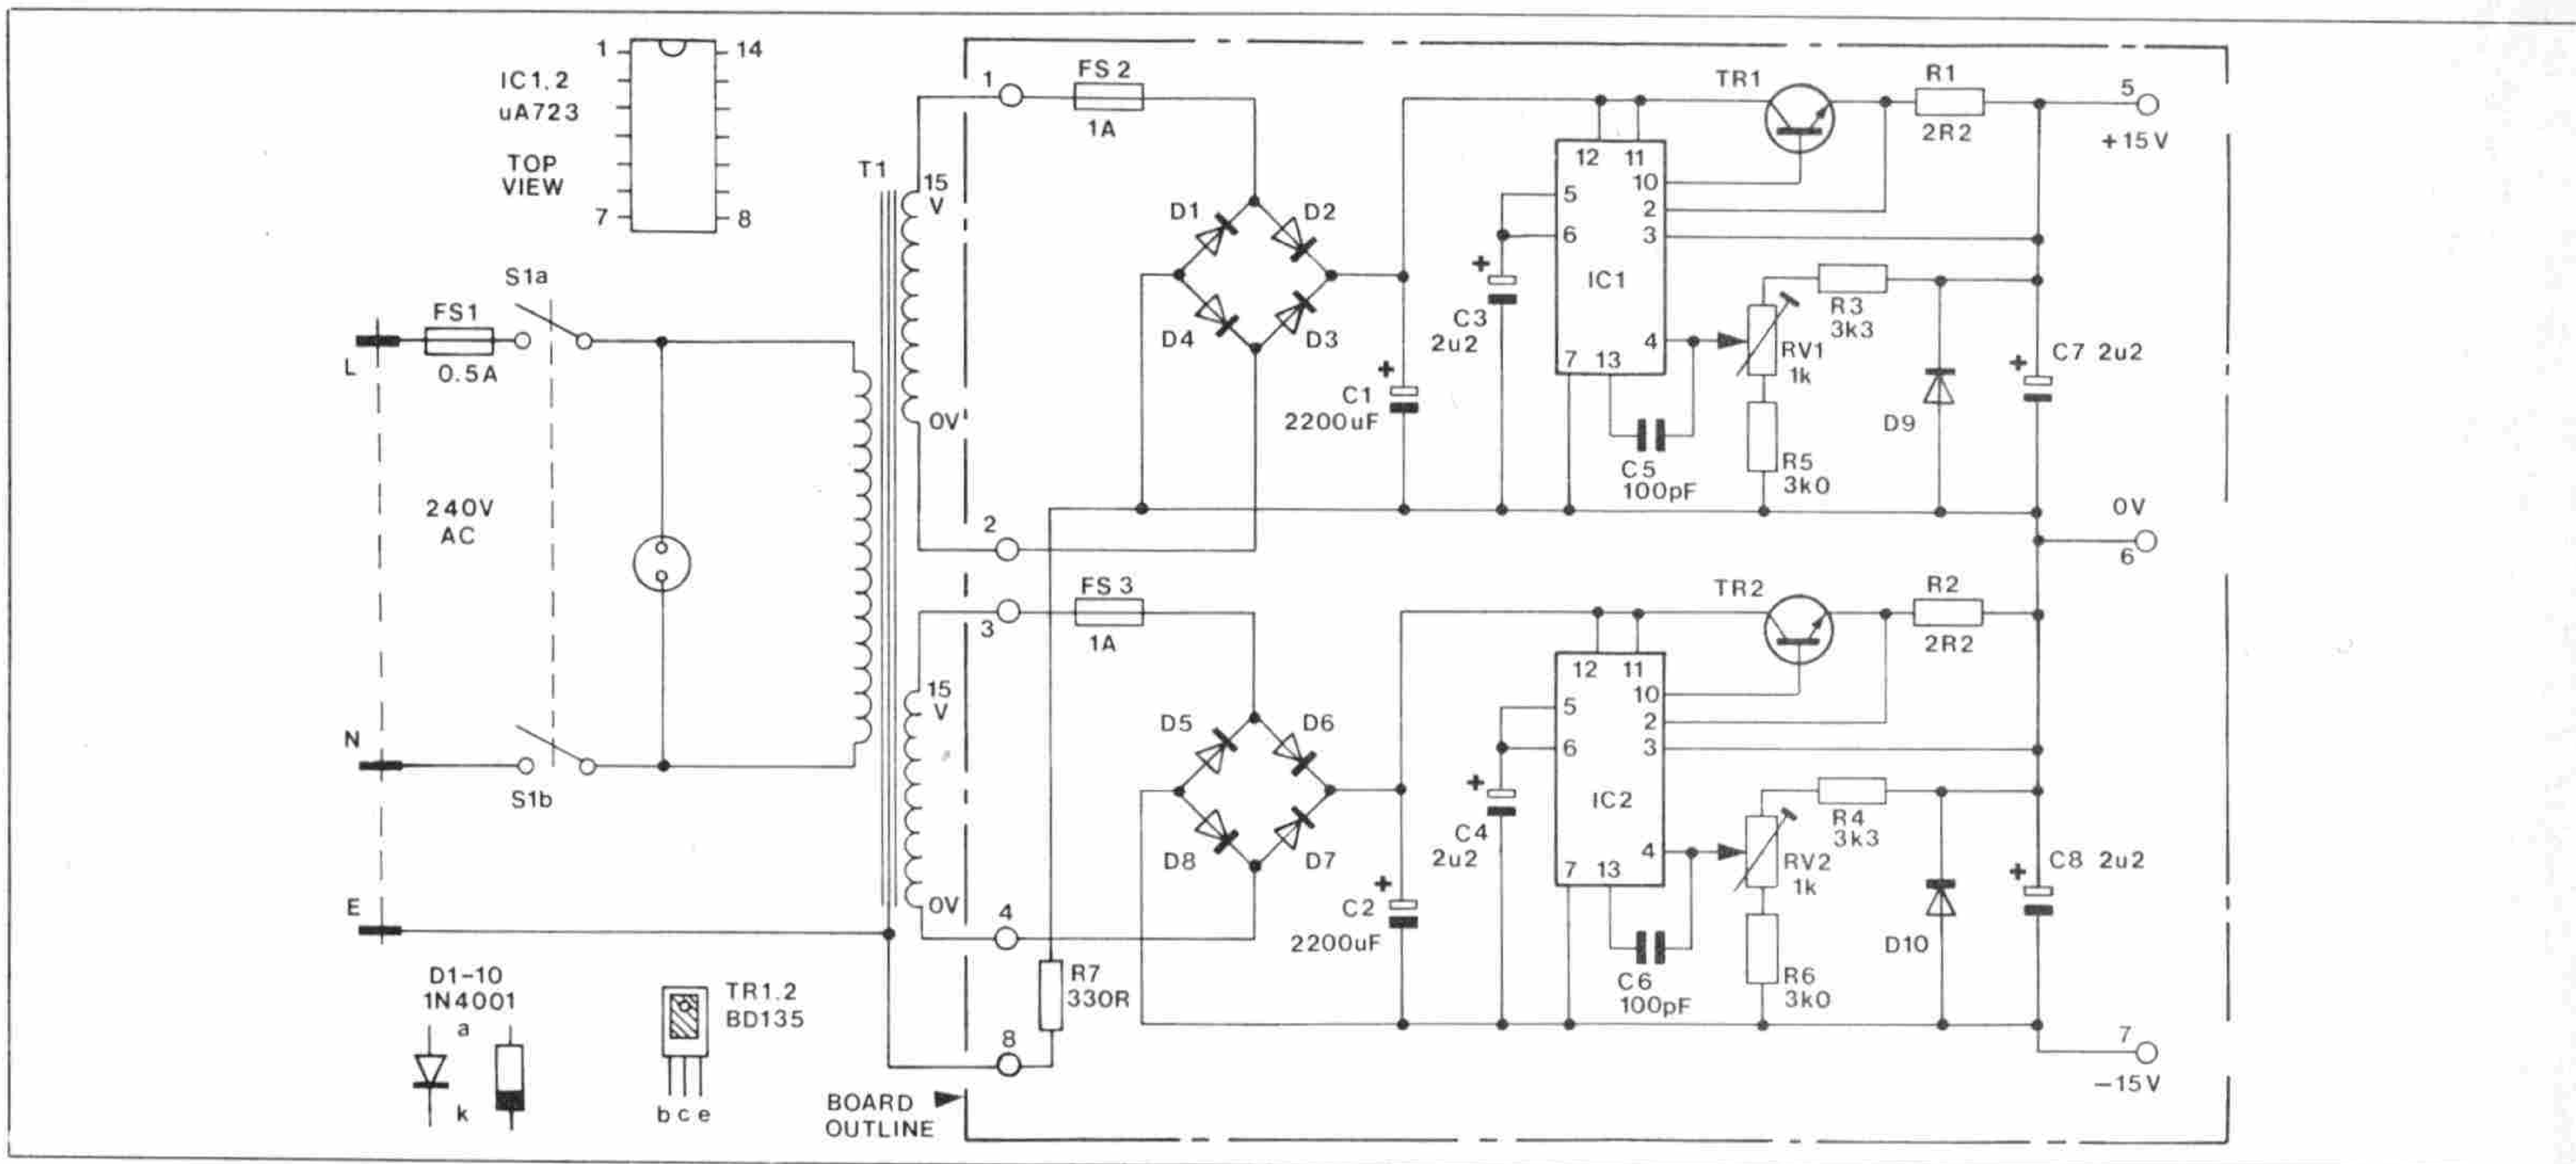

Figure 6. Circuit of the Power Supply Unit.

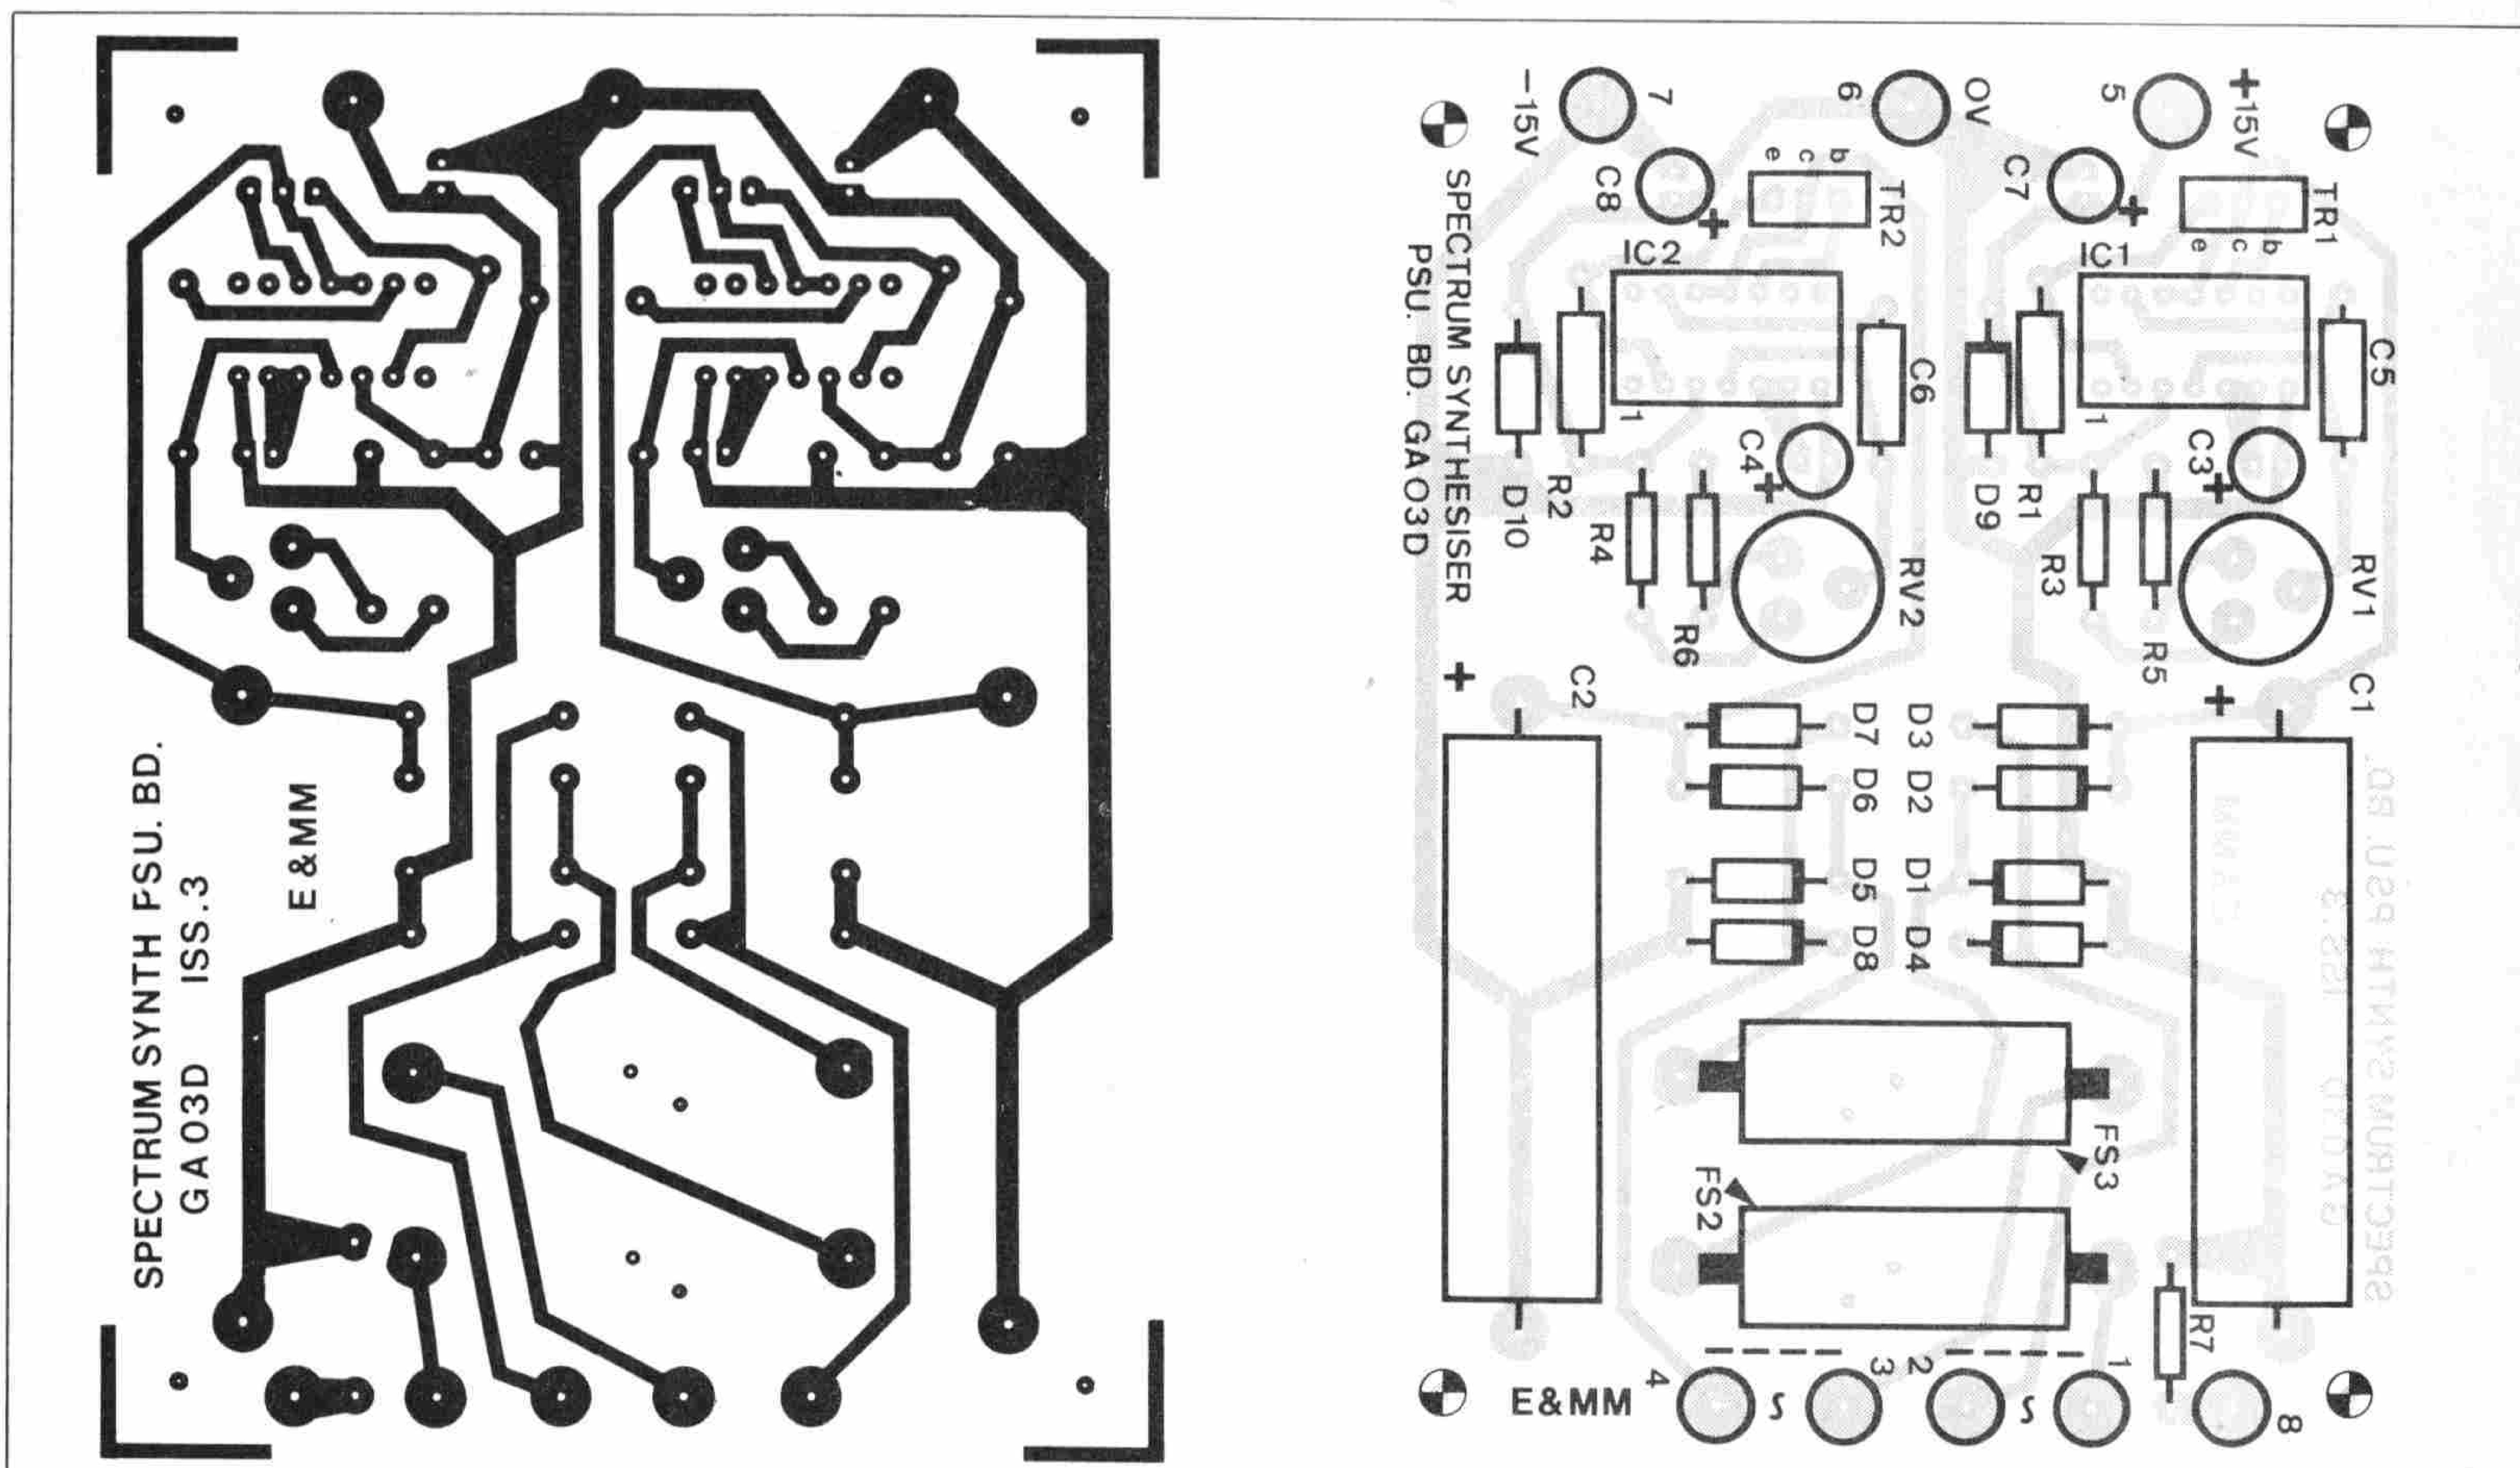

Figure 7. PSU printed circuit board.

earth connections. Strip back a length of the mains cable and cut off some of each core for connecting the switch, fuse, transformer and PCB. Connect these as shown in Figure 7, and then connect the mains cable and secure it with the grommet. The fuse is wired on the supply side of the switch, so that the switch neon will go out if the fuse blows. Fit a mains plug with a 3amp fuse and then check the wiring before going on to the next stage.

PSU Setting Up

Before inserting FS2, 3 or IC1, 2 first insert FS1, plug in, switch on, and measure the voltage across each secondary. This should be around 15V RMS. Now fit FS2 and 3 with the power off, switch on, and measure the voltage across C1 and 2, each of which should be about 21V. Switch off again, insert IC1 and 2, switch on, and measure the output voltages (across points 6,5 and 7,6). If you are confident in your constructional abilities, you can leave these checks and try the PSU with all the fuses and IC's in first time. Either way, the rail voltages should be measured and RV1 (+ve) and RV2 (-ve) adjusted so that they are exactly 15V. An oscilloscope will probably be more accurate for this than a cheap multimeter, though use a digital multimeter or a good mechanical meter if one is available.

Digisound Ltd. will be offering a full set of the CEM IC's used in this project at the special reduced price of £29.00 inc. VAT and postage. Ready-made metalwork and PCB's will be obtainable from Maplin Electronic Supplies. A cassette demonstrating the Spectrum's facilities will be made available from E&MM.

E&MM

SPECTRUM SYNTHESISER

Part 2 of this constructional series describes the keyboard controller

by Chris Jordan

Figure 2 shows the circuit diagram of the keyboard controller. Connections 1 and 2 are the bottom and top respectively of the keyboard divider chain. This is arranged in the feedback loop of IC3a, which drives a current of about 1.8mA through the divider chain. This generates 8.3V across each divider chain resistor, corresponding to a semitone, and 1V across each group of twelve, corresponding to an octave. R58 and R59 drop 1.7V so the range of key voltage is 1.7 (top C) to 5.7V (bottom C). R57 and RV3 determine the current, RV3 allowing it to be trimmed for exactly 1V/octave.

IC3b generates a signal that is used, after processing, to gate the envelope generators and key voltage sample - and - hold. With no keys depressed, the non-inverting input is held low by R60 and since the inverting input is at +0.83V (determined by R58 and R60) IC3b's output is at its negative extreme, almost -15V. When a key is depressed, the voltage at the inverting input rises to between 1.7 and 5.7V since the gate bus-bar is connected to the divider chain by the contact of the depressed key, and the output of IC3b goes high.

TR3 is an FET which acts as a voltage controlled switch in the sample-and-hold circuit around C11. It is normally held off by the negative output voltage of IC3b, via R62 and D14, but upon this going positive it is turned on and C11 charges to the voltage on the S/H bus-bar (connection point 3). Since the contact spring makes with this before the gate bus-bar, the new key voltage is always ready for sampling by the time the FET is turned on. IC5a is an FET input op-amp with a very low input bias current. This ensures that when the key is released and TR3 turns off the charge on C11 is retained with the minimum of droop. With the 50pA worst case input bias current of the buffer amplifier, it takes about 13 minutes for the pitch of the oscillators controlled by the keyboard to drop one semitone.

With the output of IC3b low, C10 is kept charged by D11, but when a key is depressed it is allowed to discharge through R65 and R66. It takes approximately 2mS for the voltage to reach the threshold of the schmitt NAND-gate IC6a, the output of which then goes low. Since D11 charges C10 very fast upon the comparator output going low, it must remain high for at least 2mS for the gate signal to be passed on to IC6b. This ensures that the effect of contact bounce upon key depression or release is eliminated and cannot cause false triggering of the envelope generators.

The external gate signal is inverted by TR4 and NAND-ed with the output of IC6a to give the key gate signal which is sent to the EG's.

If a new note is played on the keyboard before the previous one is released, a new CV is generated, but since the key gate signal remains high, the EG's will not restart their envelopes. This can be a problem when percussive envelopes are used, fast keyboard runs giving missed notes. The problem is eliminated by detecting a change in CV at the sample and hold output, and generating a key retrigger signal for the EG's. IC4a is a high-gain differentiator that produces a pulse for each change in the value of the CV. These pulses are rectified and squared up by the comparator IC4b, and lengthened by D16, R75, and C12 to a minimum of 5mS.

Contact bounce produces a very ragged CV change when a note is depressed while one is already down, and this in turn produces a multiple pulse at the output of IC4b. The circuit around IC6c generates a clean 500uS pulse from this signal — most important for external devices such as sequencers which count in response to triggers from the keyboard. When the charge on C12 reaches the threshold of IC6c, the output goes high and C14 charges via D18 and R85. After 500uS, C14 also reaches the required level, the output is forced low, and C14 begins to discharge slowly through R81. For 30mS after each pulse C14 inhibits IC6c so that no more pulses can occur at the output during this period.

Since the sample-and-hold voltage is updated before the key gate starts, a first key depression would cause an unwanted pulse on the key retrigger line. This is eliminated by D17, which holds the input of IC6d high until the gate is received.

The de-bounced gate signal from IC6a is inverted by TR5, which drives the 'key gate out' interface jack. D19 causes the gate out signal to go low in response to the key retrigger

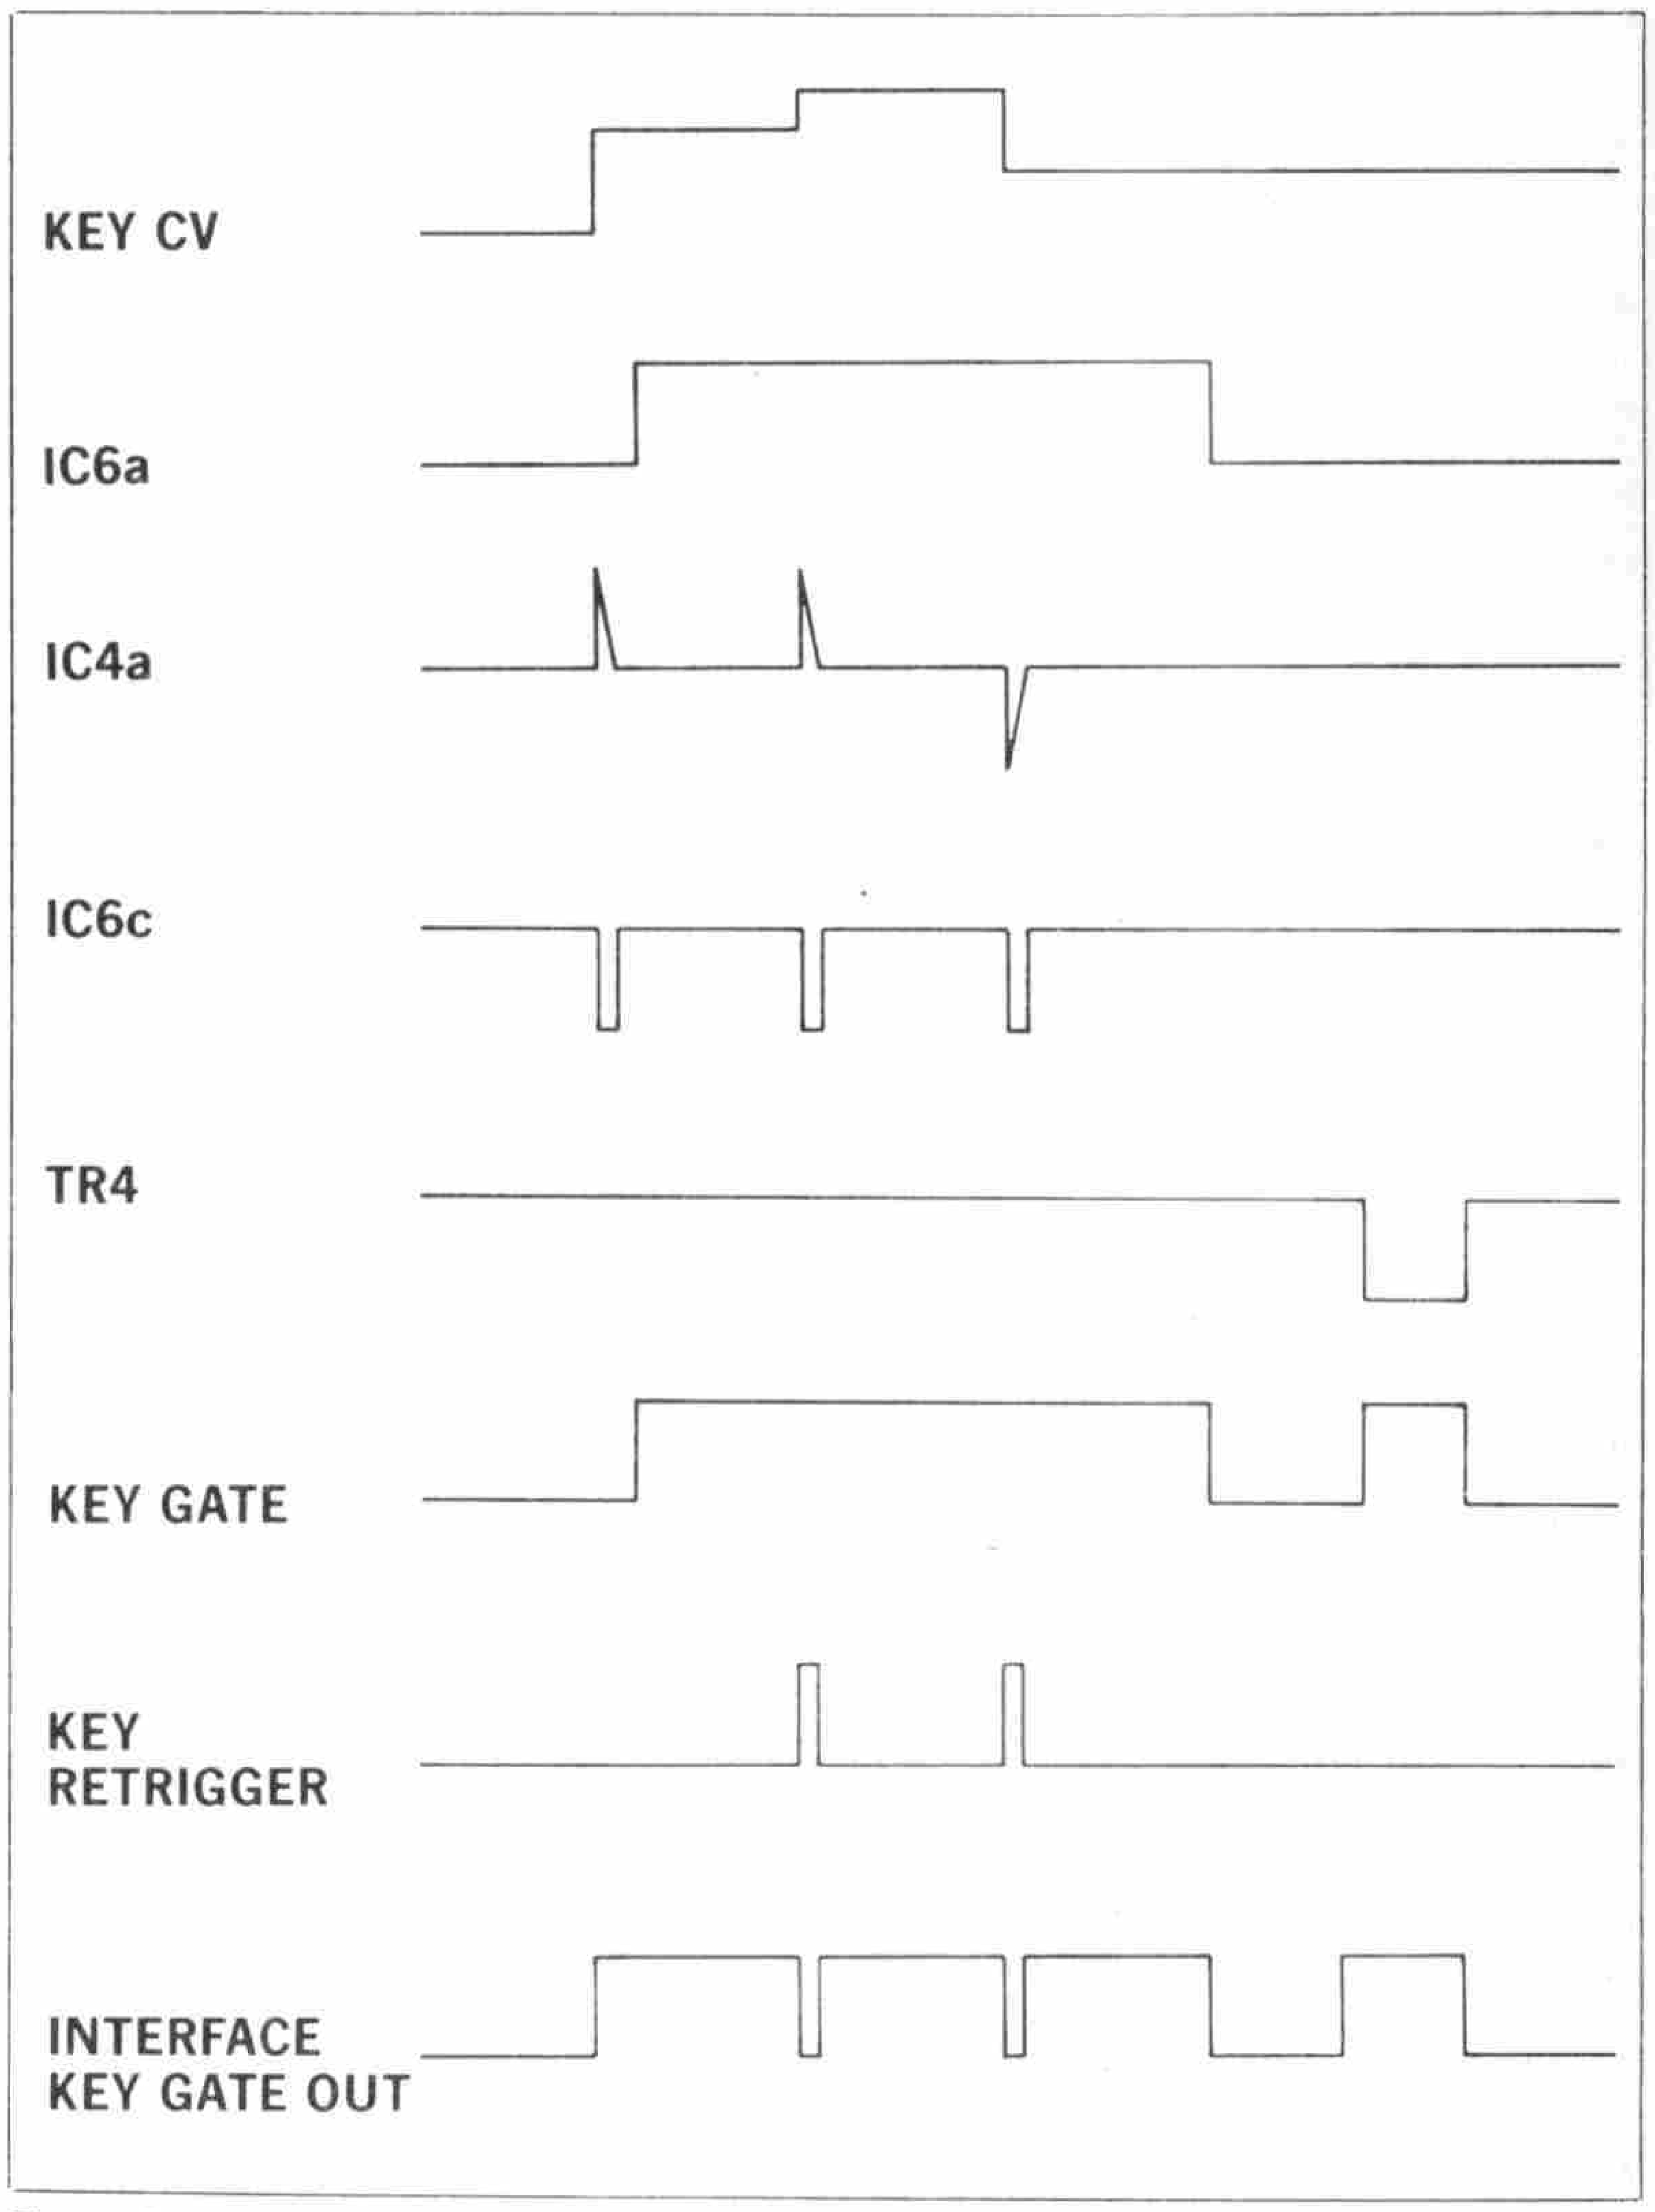

Figure 1. Keyboard controller signals.

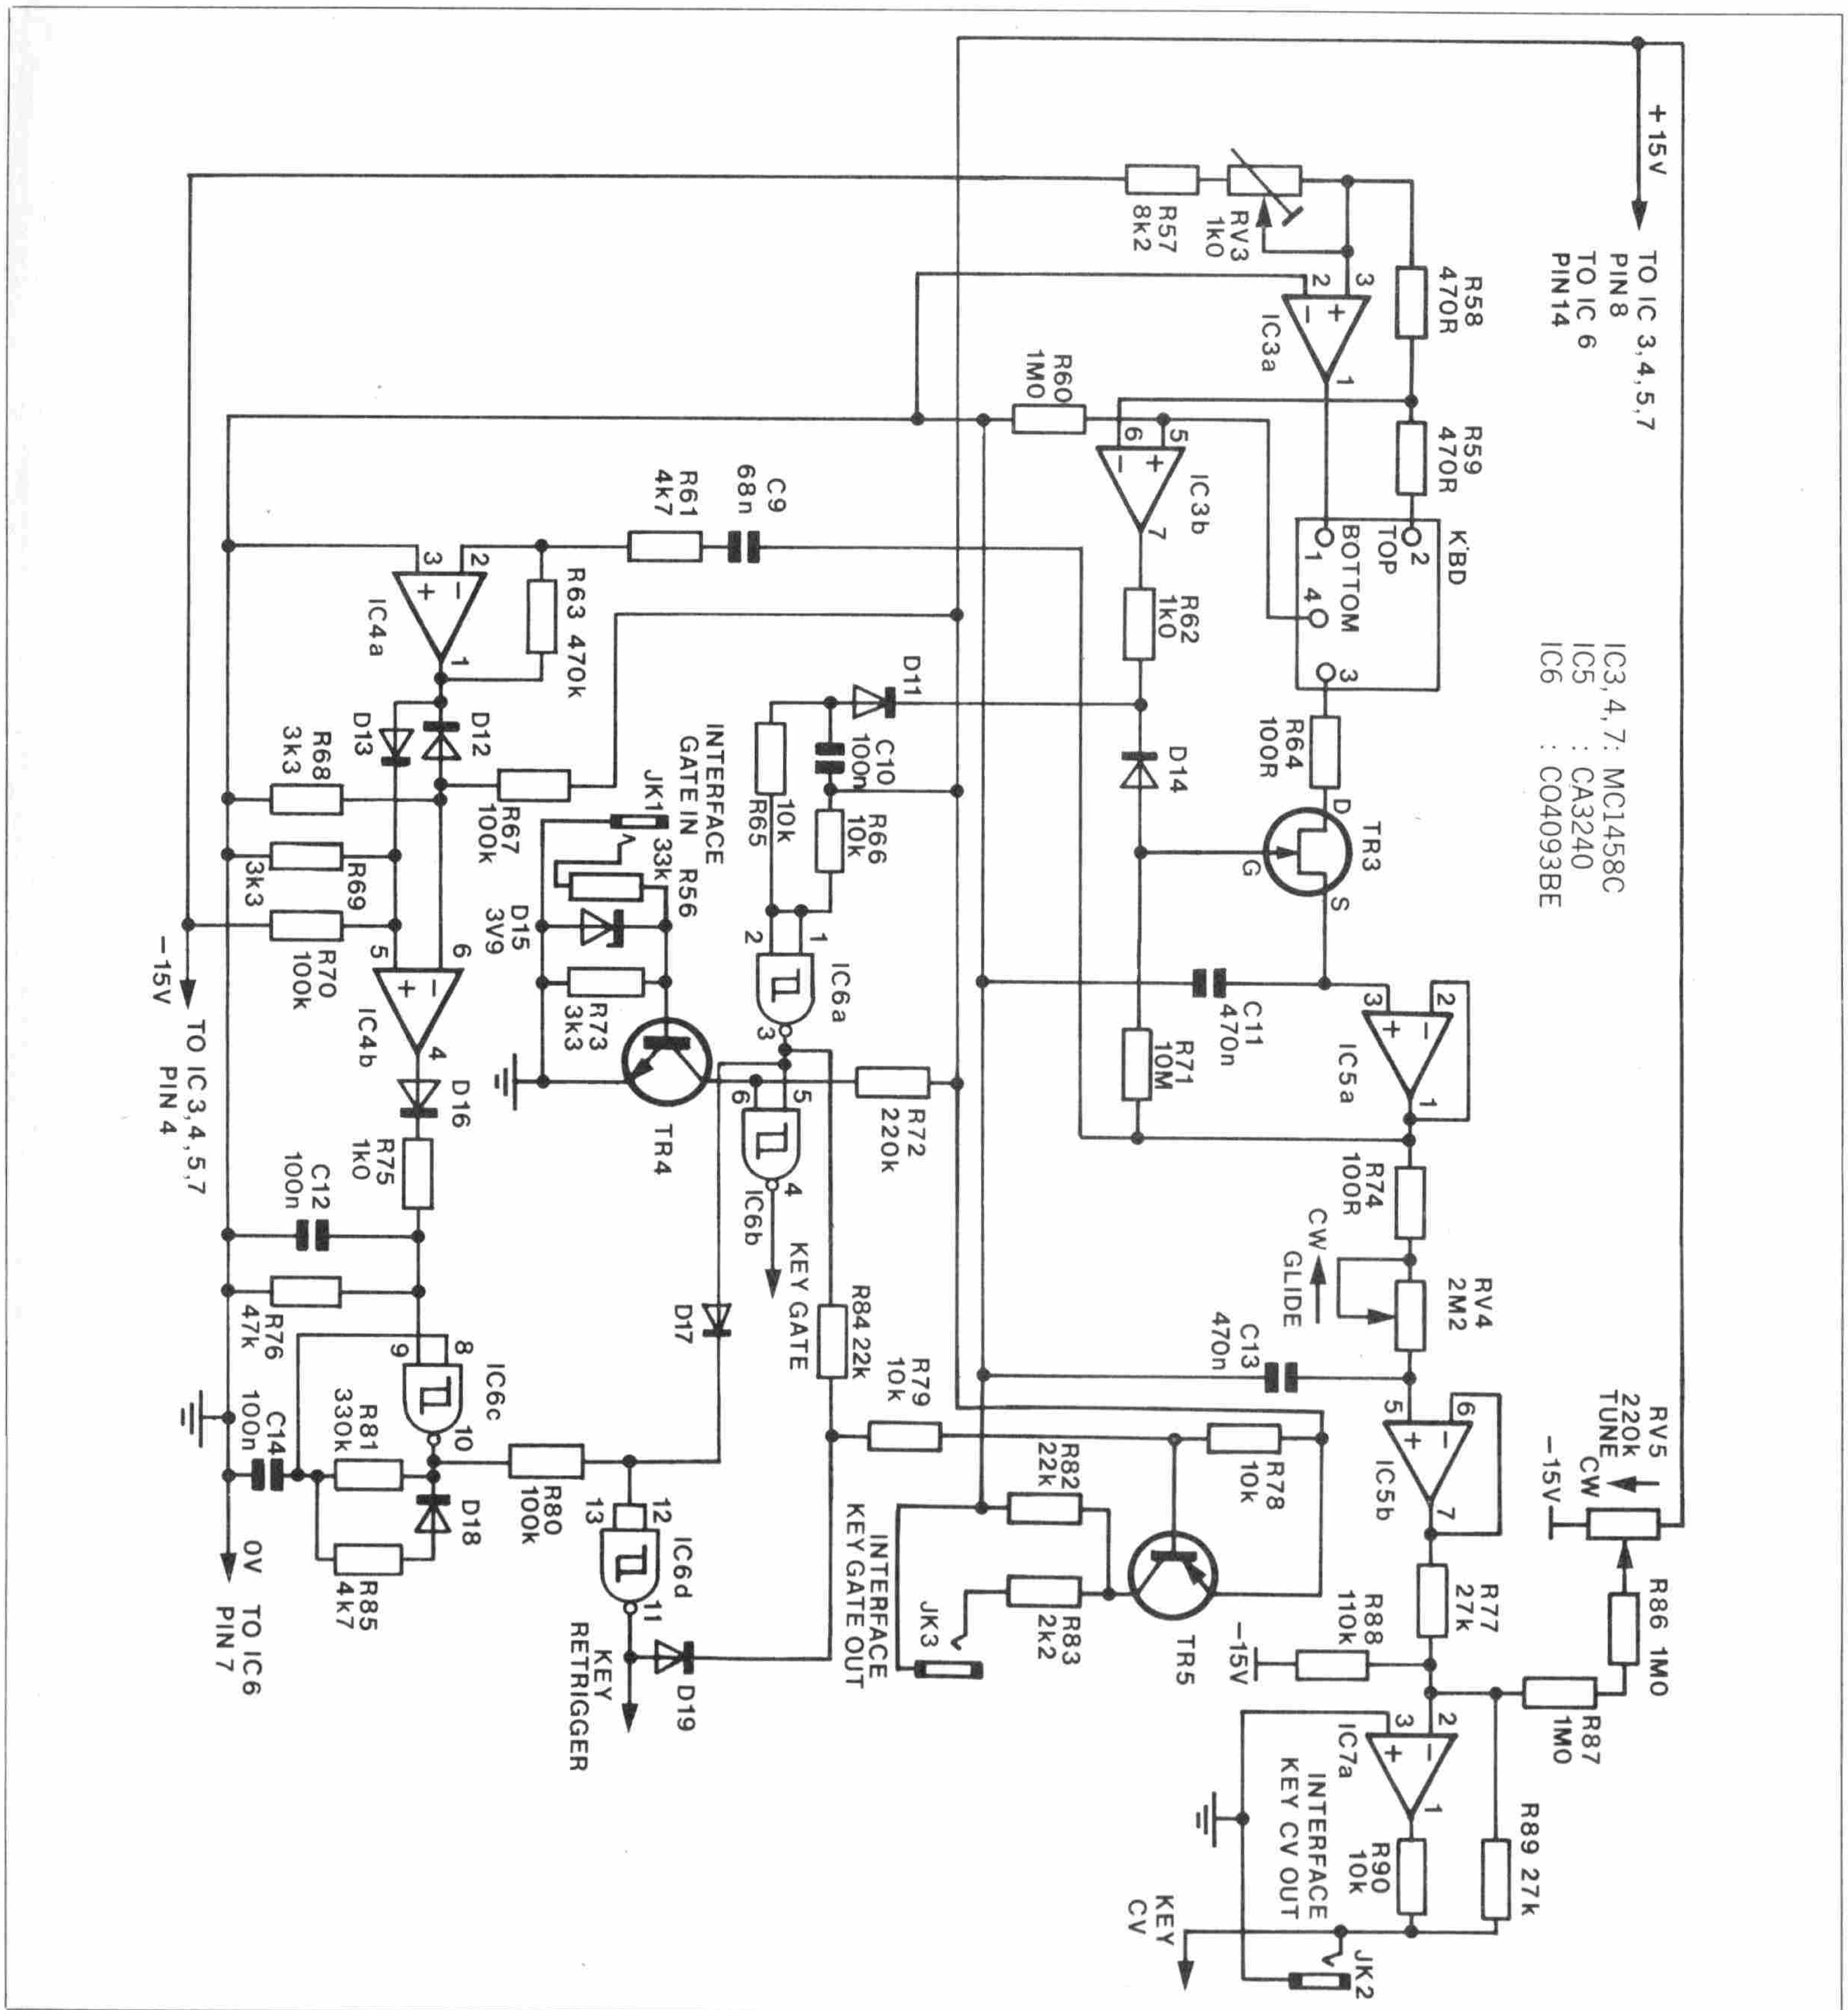

Figure 2. The circuit diagram of the Spectrum keyboard controller.

signal. TR5 is arranged to pull the output to +15V to generate the gate signal — this system allows gates from different sources to be connected together, providing an OR-function that gates the controlled device if any source signal is high.

The output of the sample-and-hold circuit (TR3, C11, IC5a) is passed to the glide circuit (R74, RV4, C13, IC5b) which produces sweeps between successive notes. The time taken for a new note voltage to be reached is controllable from almost instantaneous to five seconds for one octave by RV4. IC5b is a low input bias current op-amp, avoiding any voltage drop across RV4 that would cause a perceptable pitch error with maximum glide.

IC7a inverts the output of the glide circuit, and applies an offset so that the middle 'C' of the keyboard generates a key CV of 0V. This simplifies interfacing with additional equipment. The 'Tune' pot. (RV5) shifts the pitch up to ± 2 semitones. R90 limits the current supplied by IC7a but does not affect the voltage under normal conditions. This is required since the CV is momentarily shorted to earth when the other end of the patch lead from JK2, the 'key CV out' interface jack, is plugged into another piece of equipment.

E&MM

SPECTRUM SYNTHESISER

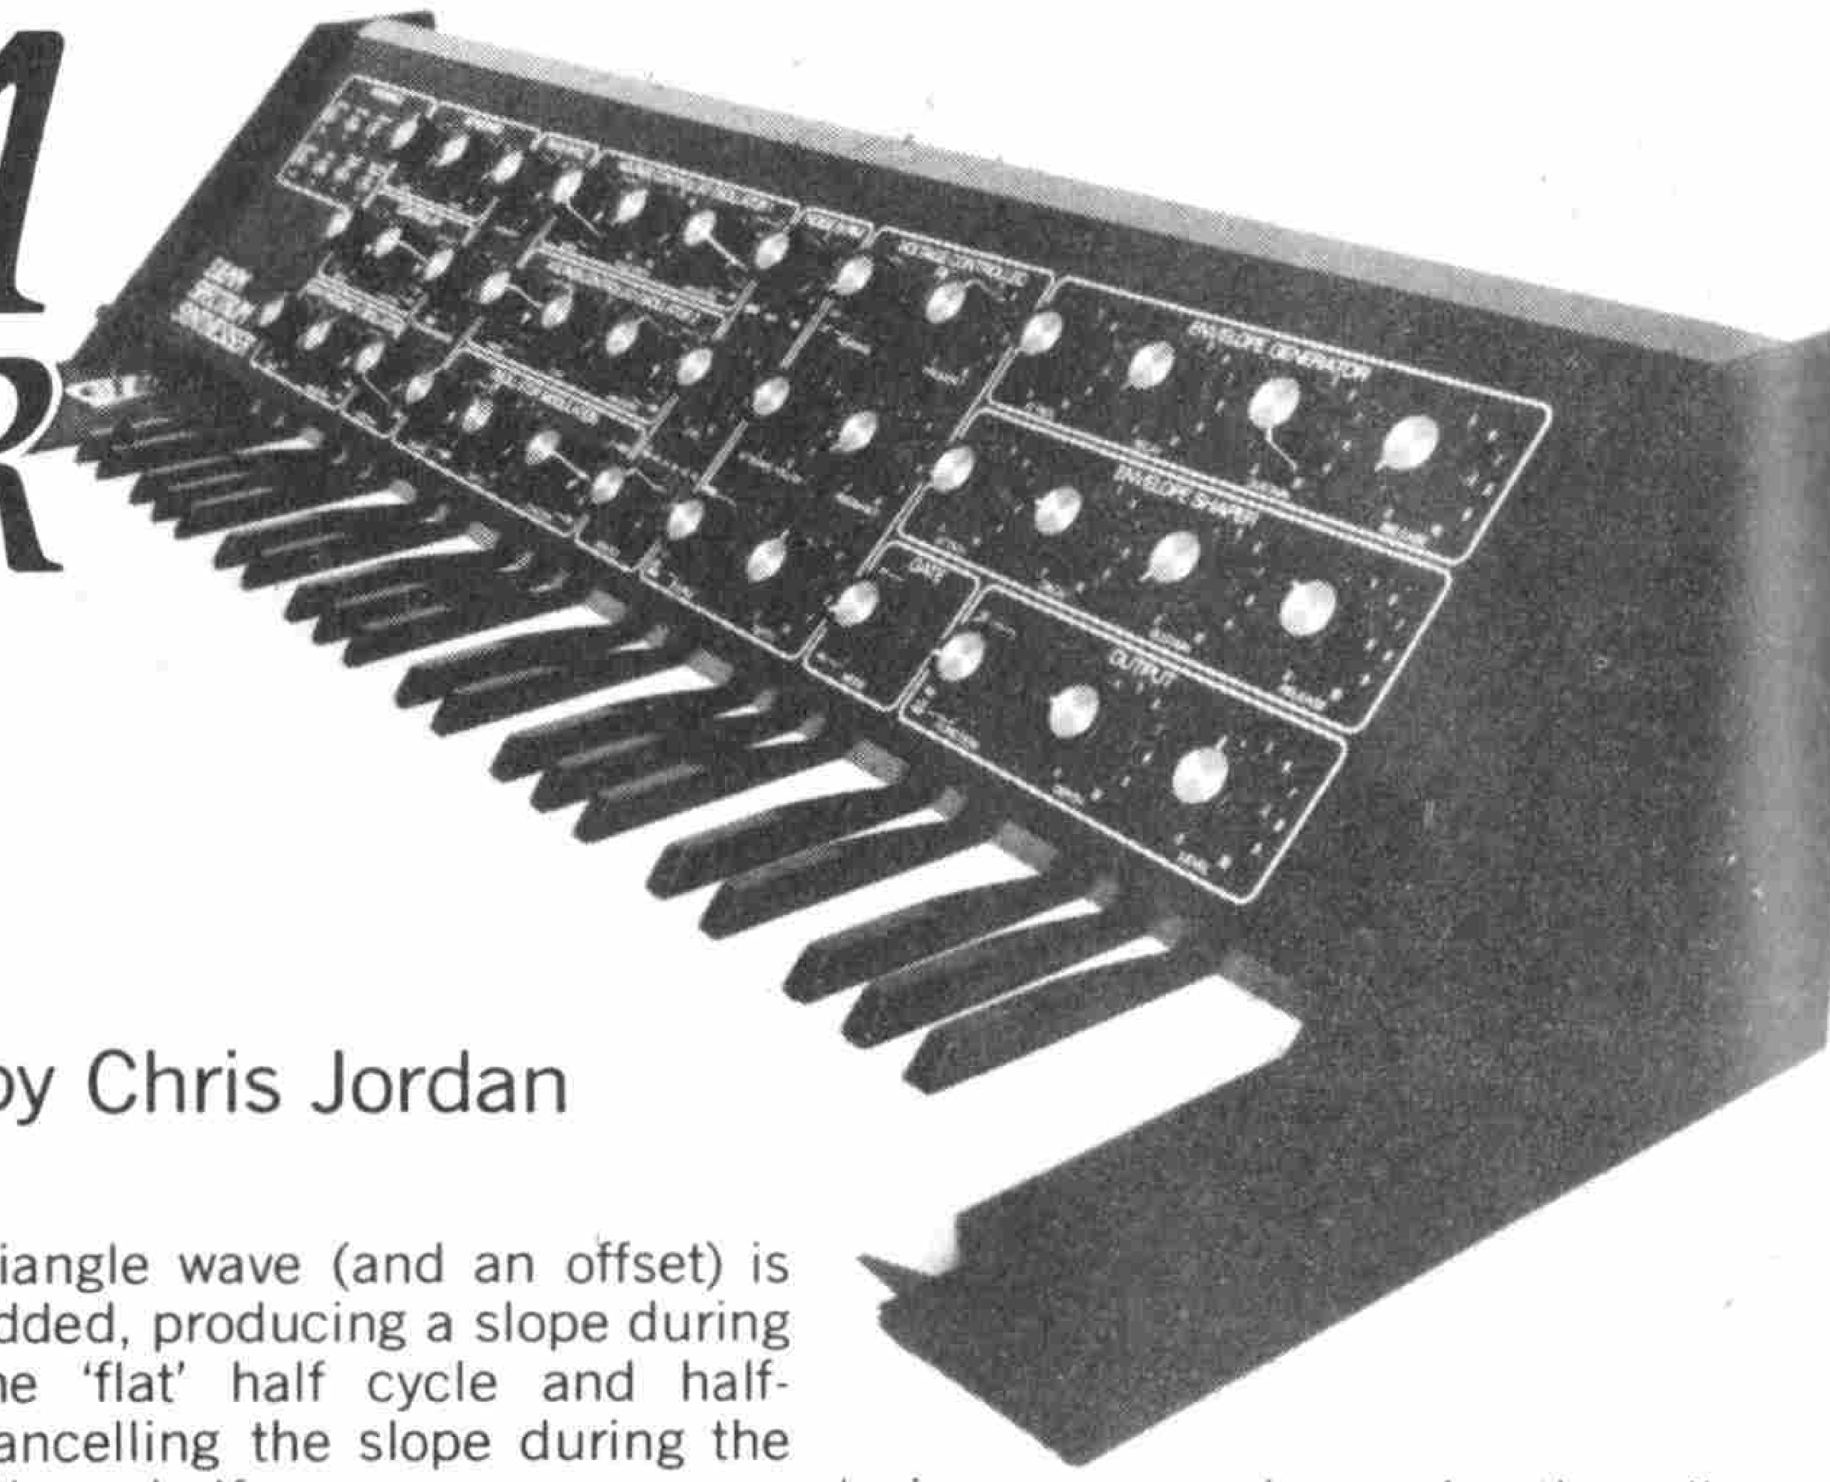

Part 3 of this constructional series describes the Low Frequency Oscillator and the Voltage Controlled Oscillator section.

by Chris Jordan

Low Frequency Oscillator

The Low Frequency Oscillator (LFO) of a synthesiser provides periodic waveforms for the control of other modules to produce modulation of pitch, timbre, amplitude etc. When the synthesiser is being used other than for simple melodic playing, the LFO is often the main control source, and must have a wide frequency range and a choice of precise waveforms. The Spectrum LFO has a range of over 1000:1, from 0.04Hz (25 seconds per cycle) to about 42Hz. Sine, triangle, ramp, and square waveforms are available, plus two additional step-type waveforms, one giving a new random voltage on each cycle, the other producing a wide range of repeating sequences. A green LED flashes to indicate the LFO cycle and is very useful for quickly checking or setting the rate. Particular attention has been paid to waveform precision, and good symmetry is retained over the frequency range. Unlike many other designs, no setting up is required.

Circuit

Figure 1 shows the circuit of the LFO. It is based around IC8, IC9a, TR8 and TR9, which form a precision triangle and square wave generator. IC8 is an integrator driven by the voltage at the wiper of RV6, the Rate control.

A low input bias current op-amp must be used for IC8 to preserve waveform symmetry since a bipolar device would drain the input current significantly at low frequencies, causing differing charge and discharge rates for C16.

IC9a is a comparator which reverses the voltage at the integrator input when its output reaches thresholds set by R100,101, so the integrator output ramps up and down between fixed levels generating a triangle wave. IC9a drives TR8,9 which are configured as an additional complementary pair output stage driving the integrator and from which feedback to IC9a is obtained. Since TR8,9 invert the output of the op-amp and R101 takes the signal back to the inverting input, the feedback is positive, causing the output to be either high or low and giving the comparator hysteresis. An additional output stage is used because the maximum and minimum output voltages of the op-amp are unpredictable and rarely symmetrical. This would give unequal times for the two halves of the cycle and waveforms which were not precisely symmetrical about 0V, since the thresholds are derived from the output of the comparator circuit.

The method of producing the rampwave is rather unusual. The triangle and square waves are mixed and half-wave rectified by IC9b. Since only positive output values are allowed, the signal is 'cut off' at zero volts when the square wave is high i.e. when the triangle wave is falling. The result is a positive going half-wave rectified ramp wave, which gives a complete ramp wave when the triangle wave (and an offset) is added, producing a slope during the 'flat' half cycle and half-cancelling the slope during the other half.

The sine wave is generated by D24-27 and associated resistors. Minimum harmonic content of a sine wave used for control purposes is not as important as smoothness of the waveform — it should have no sharp changes of gradient and should slow down gradually towards the peaks. This is achieved by two parallel diode shaping networks which act on the triangle wave. As the voltage increases on positive half cycles, D25 conducts first, and then D26 conducts just before the peak, with D24,27 acting on negative halves. The sine wave is produced by mixing the two components by R126,128 at IC10b. S2b selects the output waveform, with IC10b and its input resistors mixing the components for the ramp and sine waves and ensuring that all waveforms have the same level. The output of IC10b is the '+LFO' signal, and this is inverted by IC12a to give '-LFO'.

The 'LFO MAN' output gives the selected waveform at a level controlled by the jostick y-axis. RV7 is the joystick pot, acting as potential divider fed by +LFO and buffered by IC12b. Normal pots have a low resistance remaining between the wiper and the track connection when at the end of their travel — this 'end resistance' would leave a small signal at LFO MAN when the joystick was 'off', and this could be a nuisance when using extreme modulation depths. RV8 introduces a small amount of +LFO to the inverting input of IC12b, allowing the residual signal to be cancelled out

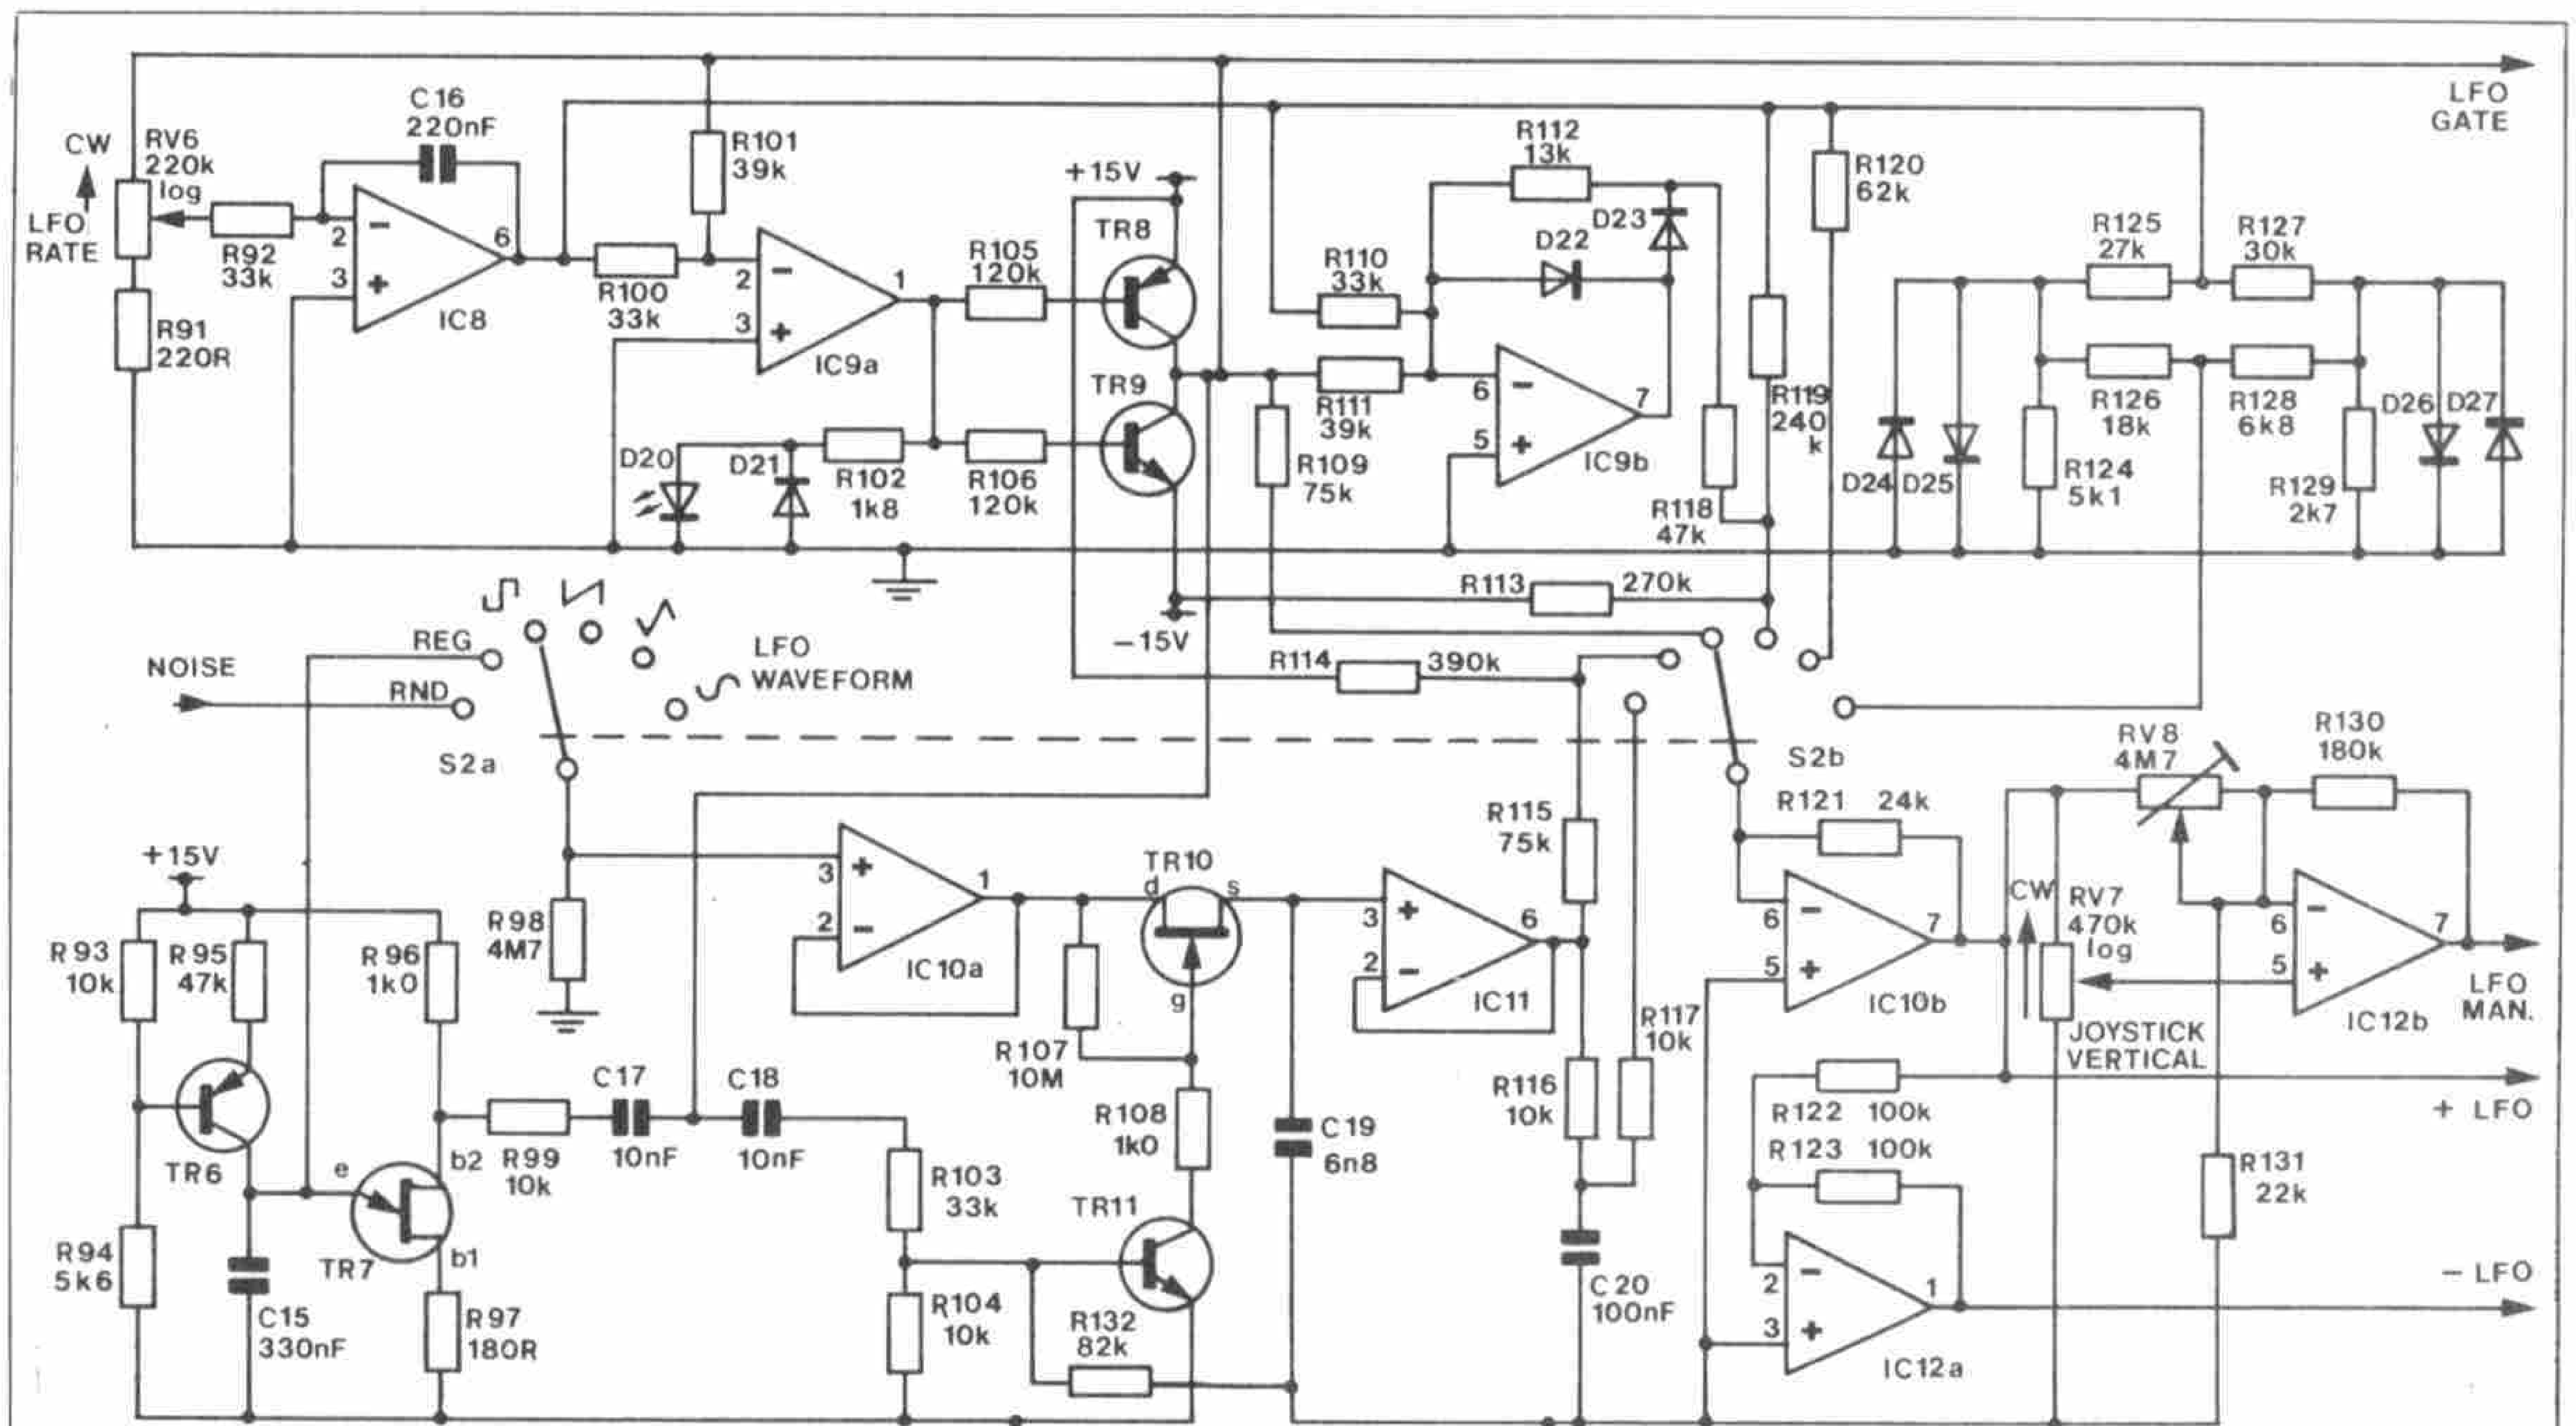

+15V TO IC8,11 Pin 7 IC9,10,12 Pin 8 -15V TO IC8-12 Pin 4	IC8 = LF351 IC9, 10, 12 = MC1458C IC11 = CA3140	TR9, 11 = BC184L TR6, 8 = BC214L TR7 = 2N2646 TR10 = 2N3819	D20 = 0.2in. LED Green D21, 22, 23, 24, 25, 26, 27 = 1N4148

Figure 1. Circuit diagram of the Low Frequency Oscillator.

and giving a very wide range of modulation depth control. This is also helped by using a log. pot. for RV7 since joysticks only move the wiper over a small section of the track and making this the 'steep' end of a log. pot. increases the effect.

The regular and random LFO waveforms are step-type functions which change level abruptly at the beginning of each cycle and remain fixed until the next cycle starts. They are produced by the sample-and-hold circuit around C19 and differ in the type of input to the sample-and-hold (S/H). The random waveform has the output of the noise generator as its source, producing a new random voltage in the range ±2.5V every cycle. The regular waveform is more complicated since the source is periodic — a 20Hz rampwave which is synchronised to the main LFO. This is generated by the oscillator around TR6,7 and C15. TR6 is a constant current souce, linearly charging C15. When the voltage on C15 reaches +5V,TR7, a unijunction transistor, turns on and discharges it to -10V via R97, from where it begins to charge again. With the regular waveform output selected, S2a connects C15 to IC10a, which buffers the ramp wave signal. TR10 is the S/H switch, normally kept off by TR11 which holds its gate negative. Upon the LFO square wave going negative at the beginning of a cycle, the pulse from C18 turns TR11 off and TR10 is allowed to conduct. C19 charges to the value of the input signal, and at the end of the sample period, which lasts about 1 ms, TR10 turns off again and the charge on C19 is held until the next sample. IC11 is FET-input op-amp connected as a voltage follower, buffering the voltage on the capacitor. A low input bias current device is necessary to minimise the drain on C19, achieving a low voltage 'droop' between successive samplings.

The output of IC11 is fed to IC10b where the S/H waveform can be selected by SW2, with the values of R114, R115 chosen to produce a ±2.5V output signal from the -10,+5V range of the sampled voltage. With SW2 in the 'Random' position, the signal is low pass filtered by R116,C20 which removes the burst of noise that appears while the sample-and-hold FET is on. Though this is only 1ms long, it could breakthrough into the audio chain when using large modulation depths of VCF cutoff frequency or VCA amplitude.

The effect of sampling a constant frequency rampwave at a regular rate is to produce complex repeating sequences of voltages, the sequence length and type being determined by the sampling and sampled frequencies. This is often used to produce note sequences by modulating a VCO with the sample-and-hold output, but suffers from the disadvantage that the slightest change in sampling frequency or the frequency of the sampled waveform changes the effect. In practice it is very difficult to get a precisely repeating sequence, rather than one which has a repetitive 'theme' that steadily changes as a part of a truly repeating sequence with a much longer period. In other words, the results are often too complex and uncontrollable to be useful, and some method is needed to restrict the S/H waveform to shorter repeating sequences. The Spectrum is unique in providing this, and does so by prematurely resetting the rampwave oscillator if it is near the end of its cycle when sampling occurs. Referring back to the LFO circuit diagram, this is achieved by C17,R99 which couple pulses from the LFO square wave to base 1 of TR7, the unijunction transistor in the rampwave generator. When the square wave goes low, the reset threshold of TR7 is effectively reduced by about 1 volt, so if the voltage on C15 is above +4V at this instant, the ramp wave is reset early and the sample-and-hold receives the voltage at the start of the next ramp cycle, i.e. -10V. The rampwave generator then runs normally until the next time it falls above +4V on a sample, whereupon it is reset and the sequence is repeated exactly. The time taken for this to occur depends upon the frequency ratio, but since the synchronisation is quite weak, sequences from very short to quite long are easily obtained and very long sequences are terminated when the premature reset condition arises.

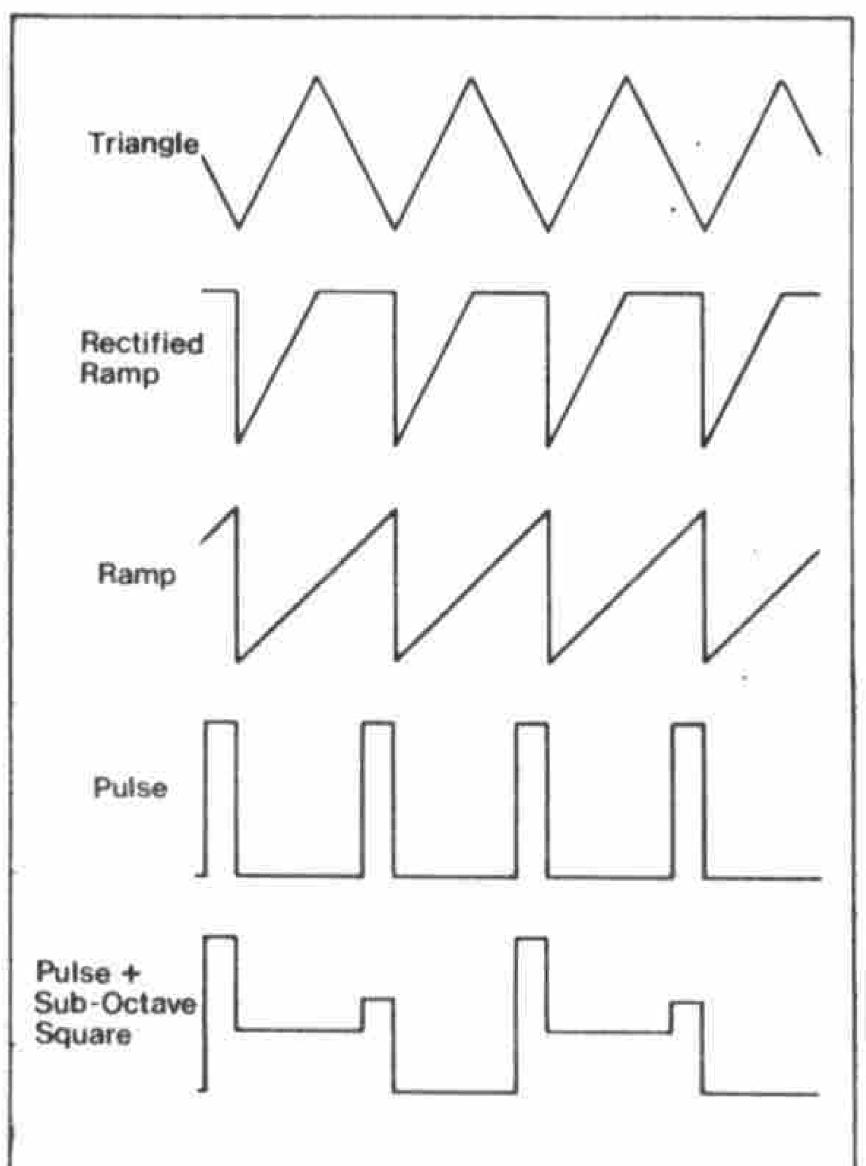

Figure 2. Basic VCO waveforms.

The LFO square wave is sent to the envelope generator and shaper separately from the waveform selector switch and modulation routing, where it can be used to gate the envelopes repeatedly.

VCO's and Associated Circuitry

The Voltage Controlled Oscillators (VCO's) are the heart of the analogue synthesiser, and to a great extent determine the overall quality of the instrument. In exponential synthesisers they must be carefully designed to give an accurate and temperature compensated control scale; this normally makes them the most expensive sections and requires complex setting up.

In most small synthesisers the Voltage Controlled Filter is the primary timbre-determining section, with variations between designs responsible for the characteristic sounds of different instruments. The VCO's play a lesser part in tone forming, with a limited choice of basic waveforms available to the player. The Spectrum synthesiser incorporates design techniques never before used in an instrument of this type to provide a very wide range of different timbres from the oscillator section by using the two VCO's in combination.

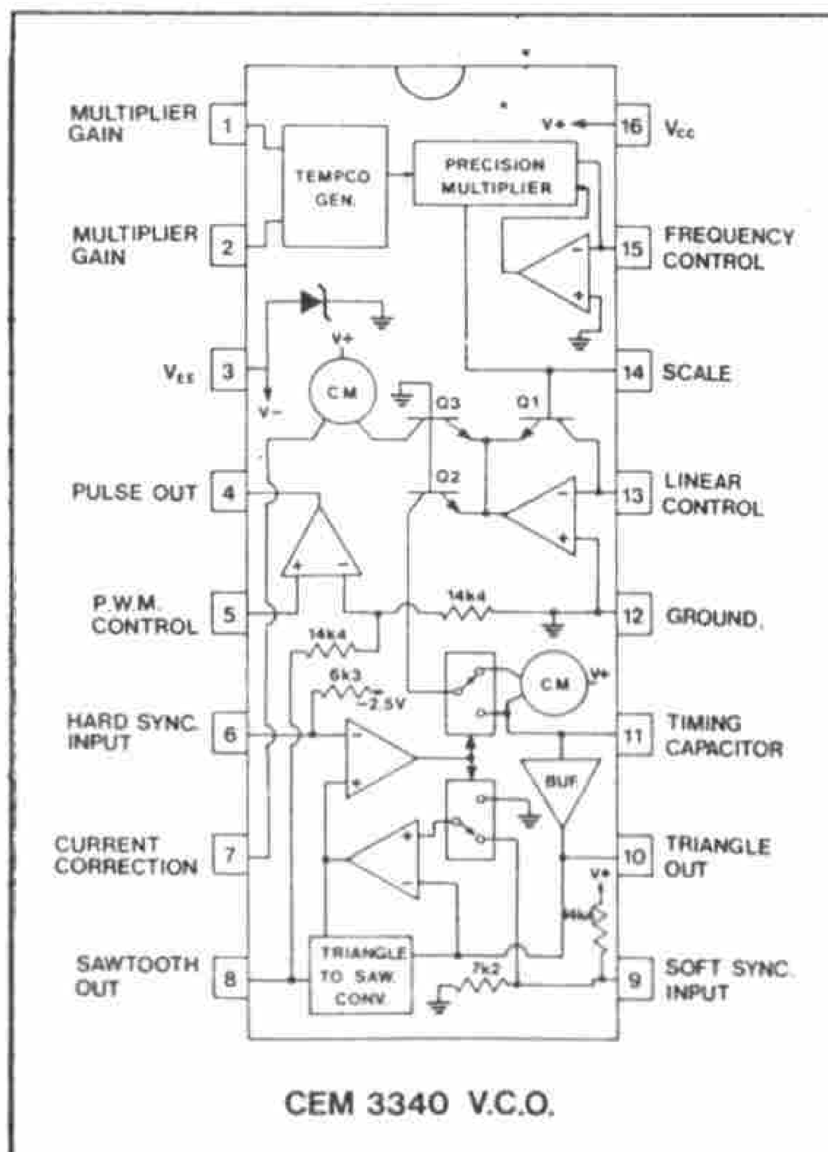

Figure 3.

A unique feature is the ability of one oscillator to sound at harmonics of the other oscillator only — when used with the regular and random LFO waveforms this provides sequencer effects that can be transposed from the keyboard. This and other applications present new musical possibilities not previously open to players of low-cost synthesisers.

Circuit

Figure 4 shows the circuit diagram of the Voltage Control Oscillators (VCO's) and noise generator, which together form the source of audio signals for the synthesis of all available sounds. The oscillator control circuitry and the sections that combine the VCO signals by frequency modulation, synchronisation, and ring modulation are also included.

Each VCO uses the CEM 3340 IC, which is specifically designed for this kind of application, allowing a versatile and precise VCO to be built with great improvements in cost, component count and specification over discreet designs. The CEM 3340 was fully described by Charles Blakey in 'IC's for Electromusic', E&MM March '81, so except where its usage in this design is unusual, we shall not discuss it in great depth here. The internal diagram is shown in Figure 3. The device is an exponential VCO with linear FM, sync, and pulse width control inputs. In Figure 5, IC15 and IC16 are the basis of VCO 1 and VCO2 respectively, and pin 15 of each is the exponential control input. This is a virtual earth summing node so each of the required signals for VCO pitch control are routed to this point via a resistor whose value which determines the control relationship (the amount of pitch change for a given voltage change). With the scale trim presets correctly set, 100k gives the required keyboard control relationship of 1V/Octave. the key CV signal is fed to VCO1 and VCO2 via R162 and R163 respectively, which are 100k 1% metal film resistors with a temperature co-efficient of 100ppm/°C. The precision is not important since the scale is trimmed, but the low temperature co-efficient is required to ensure that the control relationship remains constant with varying temperature. IC15 and IC16 are internally compensated for temperature changes, but stability of external control signals is just as important where it affects control scale.

The VCO CV interface socket accepts an external voltage from a device such as a sequencer for additional precise control of the VCO's. The voltage is buffered by IC7b and fed to pin 15 of IC15 by R147,R164 and RV21, and to pin 15 of IC16 by R148,R165 and RV22. Though 100k 1% resistors would give a control scale as precise as that for the keyboard, the external CV must match key CV for scale exactly, so RV21 and RV22 are included. S5, RV15,S7, and R157-161 perform the modulation routing for the VCO's. S5

selects the source from among the envelope generator, low frequency oscillator, and noise generator, and RV15 controls the depth of modulation from 0 to 5 octaves when controlling pitch. A logarithmic pot. is used to provide fine control at low modulation depths. S7 selects the modulation function from pulse width, where at maximum depth the range is 50%, either VCO, or both VCO's simultaneously. The 'Off' position enables a modulation effect to be preset and then switched in when required.

The controller enables the joystick or an external device to control either or both oscillator pitches, pulse width, or filter cut-off frequency with variable depth. IC14a amplifies the voltage from the wiper of RV13, the x-axis joystick pot. With the controller in/out socket unused, RV14 controls the amount of joystick voltage modulating the function selected by S6. The joystick voltage is available at the controller socket for control of additional equipment, or a foot pedal wired as a variable resistance to earth can be connected to control the selected function. The joystick voltage can be overridden by patching in a voltage from an external low-impedance output. A signal from a high-impedance output will be mixed with the joystick voltage.

Each VCO has a range selector switch which transposes the pitch up or down over a total range of six octaves. The voltages for the different ranges are provided by the potential divider composed of R133-138, RV9-12. S3 and S4, the range switches, select 0V for 64', 2.5V for 32', 5.0V for 16' and so on to 12.5V for 2', the top setting. The basic pitch for 64' is set by the positions of RV17 for VCO1 and RV18 for VCO2, each of which applies a fixed control voltage to its respective VCO control input. The basic pitch for 32' is then set by RV19 for VCO1 and RV20 for VCO2, with the rest of the ranges having their own presets in the potential divider. The selected voltages are not sent directly to the VCO's but are buffered by IC13. This prevents the currents taken by R145, RV19 and R146, RV20, from affecting the voltages on the divider, which would otherwise cause the position of the range switch of one VCO to effect the pitch of the other. C27 stores the last selected voltage while S3, which must be break-before-make to avoid shorting out sections of the divider, is between switch positions. On many synthesisers, changing the oscillator range causes a spurious pitch to be generated, which often appears as an annoying 'blip' if a note is sounding. C27 maintains the pitch during the changeover and allows perfect octave switching while playing. C28 performs the same function for VCO2, and R141, R142 are included so that upon either range being changed, the charge currents of C27 and C28 are kept low enough to eliminate any perceptible momentary pitch drop due to drain on the divider.

One special feature of the CEM 3340 is the linear frequency modulation (FM) input, which allows the frequency of the VCO to be modulated by an audio frequency signal for the creation of new timbres. The current at this input (pin 13) is multiplied by the exponentiated pitch control voltage, so that a constant percentage FM depth is maintained over the range of the oscillator (see 'Advanced Music Synthesis', E&MM March '81). This is ideal for a keyboard-based synthesiser such as the Spectrum, since it allows a FM tone to be set up and played from the keyboard in the same way as a simple waveform. S8 is the FM & Sync function switch. In the 'FM' and 'FM + Sync I' positions, the triangle output of VCO1, from IC19a, is fed to the linear FM input of IC16. C36 removes the DC offset from the triangle wave and is arranged with R190 before S8 and RV28, the FM Depth control, so that the depth can be altered without the charge on C36 changing and causing a brief unwanted frequency shift. The value of R183 has been chosen to give just under ±100% frequency modulation depth with RV28 at maximum.

The CEM 3340 is equipped with synchronisation inputs which can be fed with pulses from another VCO to lock the VCO's to the same frequency. The 'hard sync' input accepts positive and negative going pulses which cause the triangle wave to reverse direction during its rising and falling sections respectively. The 'soft sync.' input gives access to the potential divider that produces the upper threshold voltage for the triangle wave, and by applying negative pulses to this

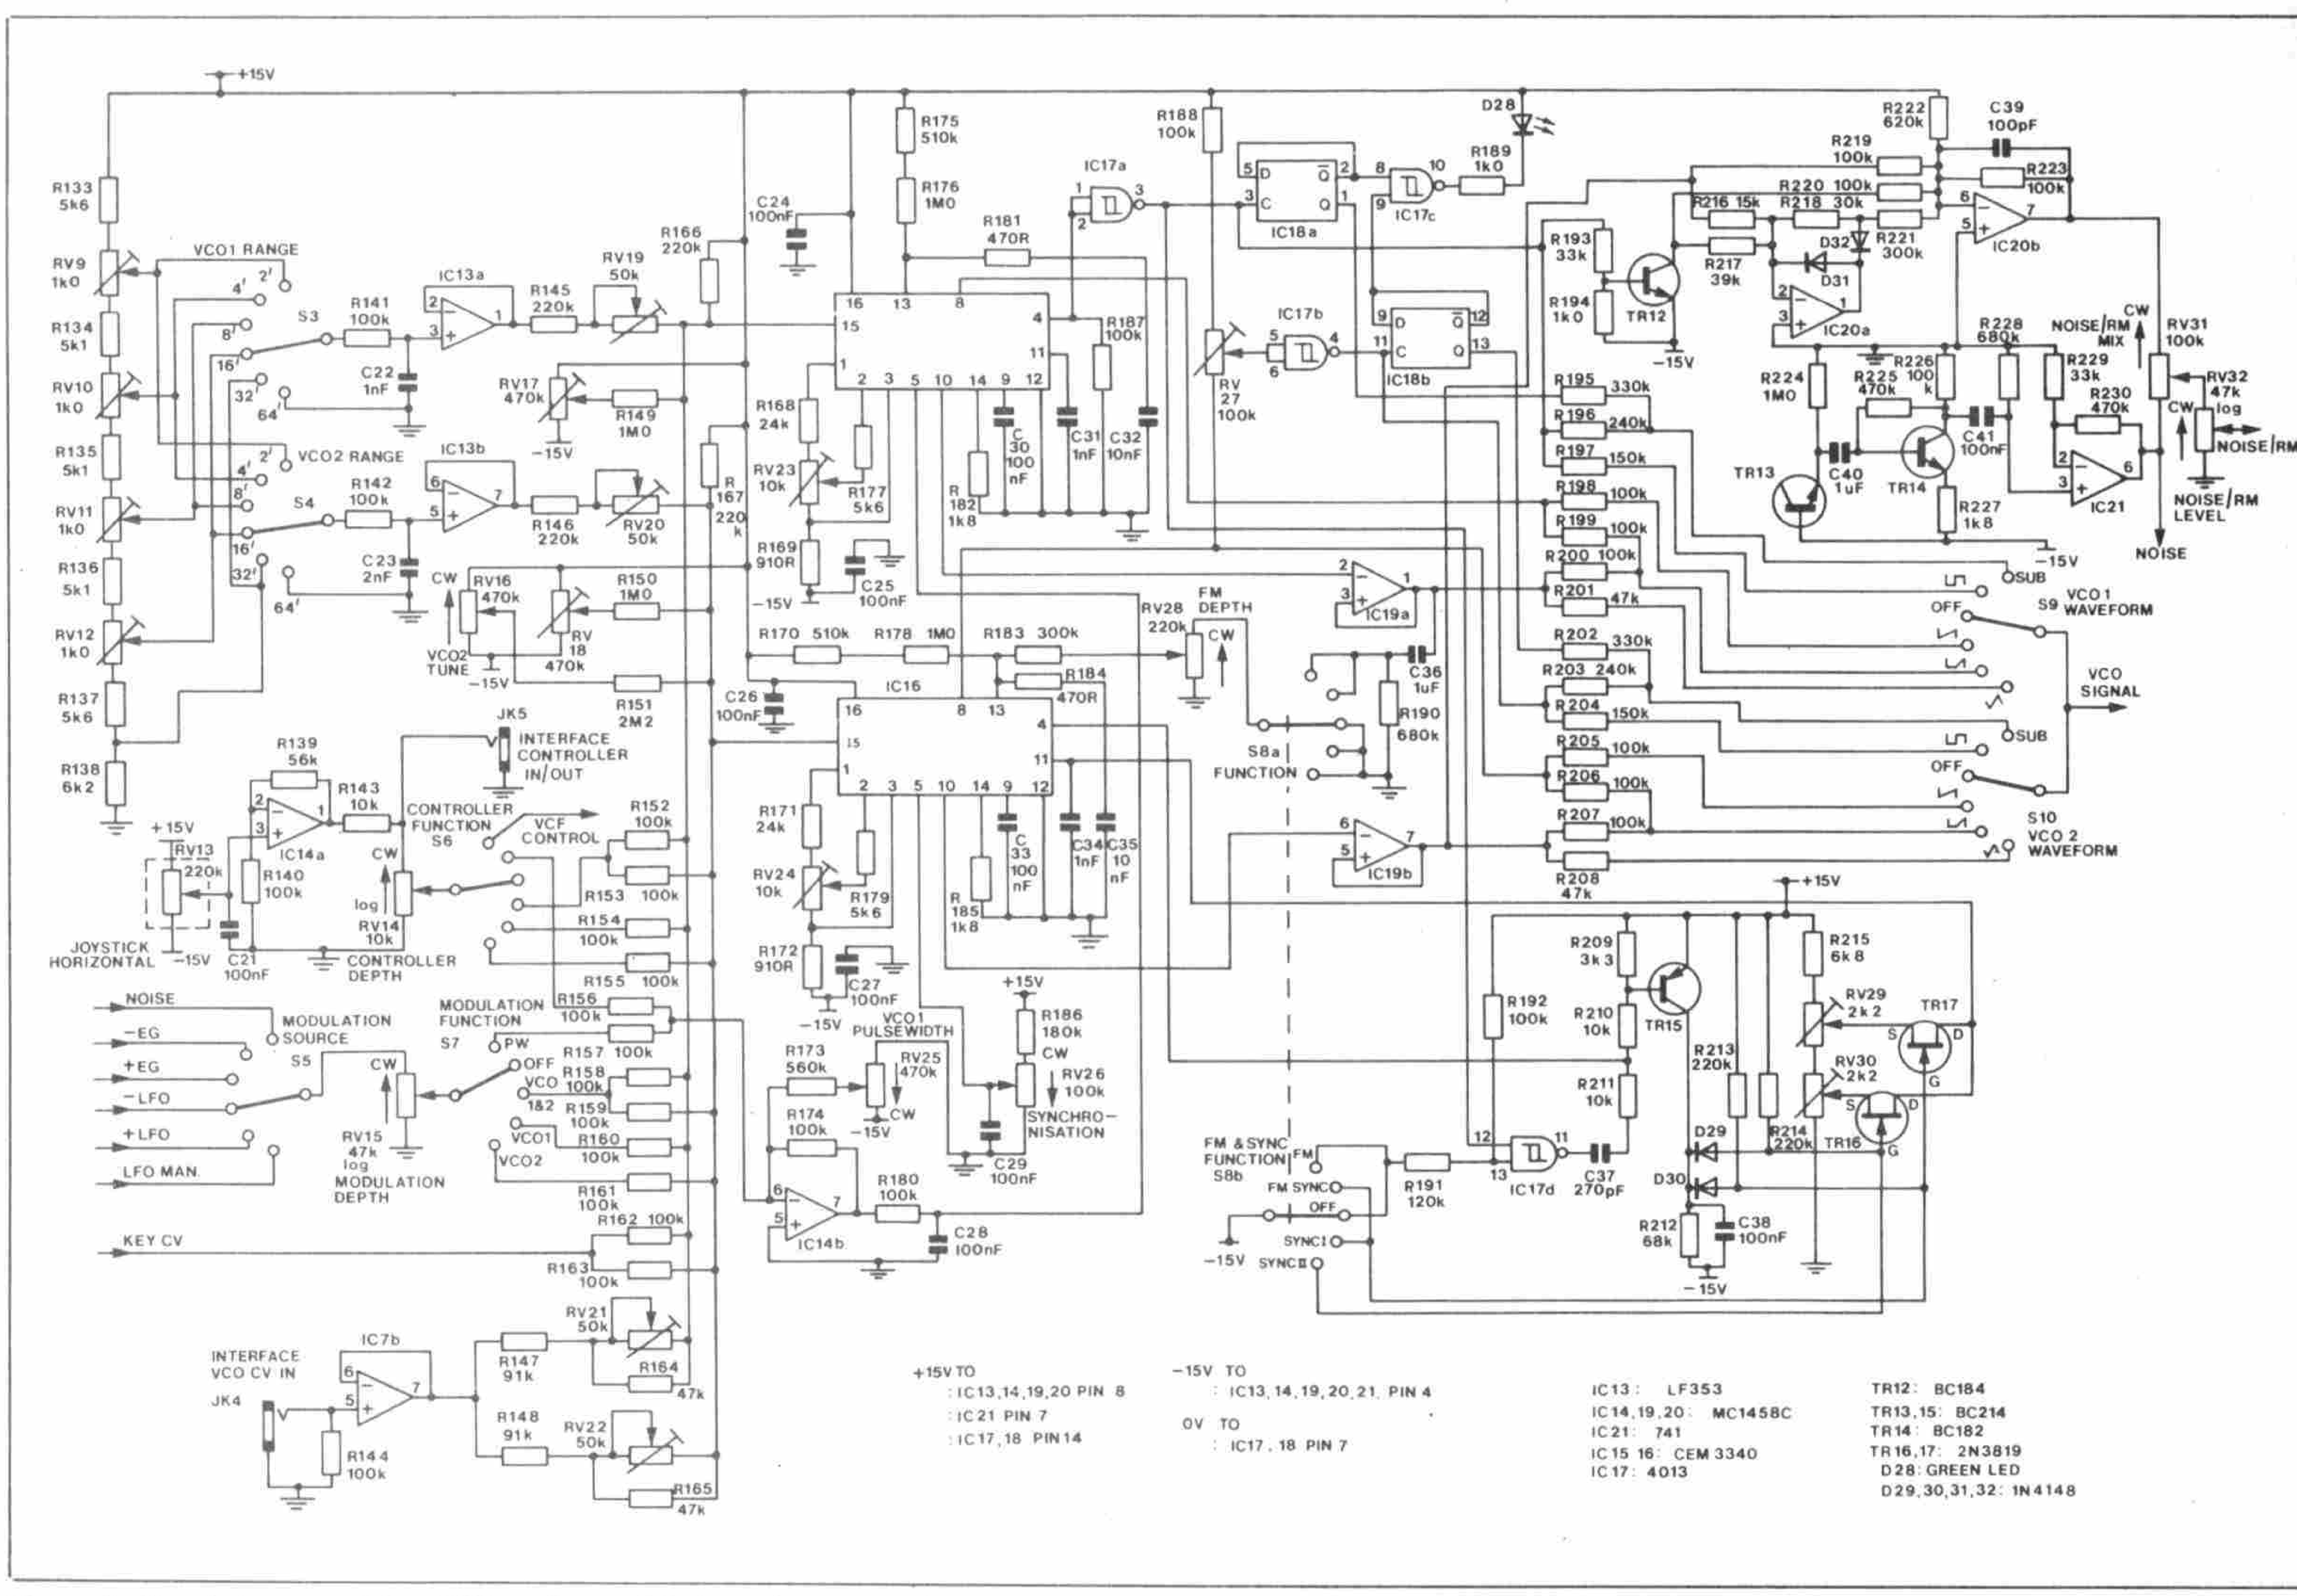

Figure 4. Circuit diagram of the VCO's and associated circuits.

point the triangle wave is reversed at its upper peak when it reaches the point at which the input pulses cause the threshold to drop below the level of the waveform. Neither of these methods provide true synchronisation since this relies on the waveform being reset to a fixed point each time, rather than merely reversing its direction. The sync inputs provided do enable the waveform to be synchronised to the frequency of the input pulses, so strictly it is correct to call the effect synchronisation, though 'hard' and 'soft' normally refer to different degrees of the same effect, with hard sync causing unconditional reset of the waveform, and soft sync causing reset if the waveform value at that time is in a particular range, usually above a certain level. The synchronisation facilities provided on the CEM 3340 are unsuitable for the creation of new waveforms, the most useful property of true sync, so the Spectrum uses additional circuitry to achieve this.

The synchronisation circuit appears in the bottom left hand corner of Figure 4. S8b is the pole of the FM & Sync Function switch that controls this circuit. When sync is off (in the 'Off' and 'FM' positions) pin 13 of IC17d is held low blocking the pulse wave from VCO1, the 'master' oscillator. When sync is selected, the pulse wave is inverted by the NAND gate and the falling edges are differentiated to give 10us wide negative pulses that turn TR15 on. TR16 and TR17 are FET's that provide a low resistance path from C34, the integrator capacitor of IC16, to the potential divider R215,RV29,RV30 when either gate is allowed to go high. Without sync selected, the FET's are held off by R212 via D29 and D30. With S8 in the 'Sync I' or 'FM + Sync I' position, the gate of TR17 is connected to -15V holding it off, but on each sync pulse R213 is allowed to turn on TR16, and C34 discharges to the voltage set by RV30. With Sync II selected TR16 is held off and TR17 discharges C34 to the voltage on the wiper of RV29. Hence at the end of each cycle of VCO1, VCO2's waveform is reset to one of two positions depending on which type of synchronisation is selected.

As can be seen from the internal diagram of the CEM 3340 the voltage on the integrator capacitor at pin 11 is buffered to drive the comparator, triangle wave output, and ramp wave shaping circuit. The comparator switches the threshold and direction of the triangle waveform when the selected

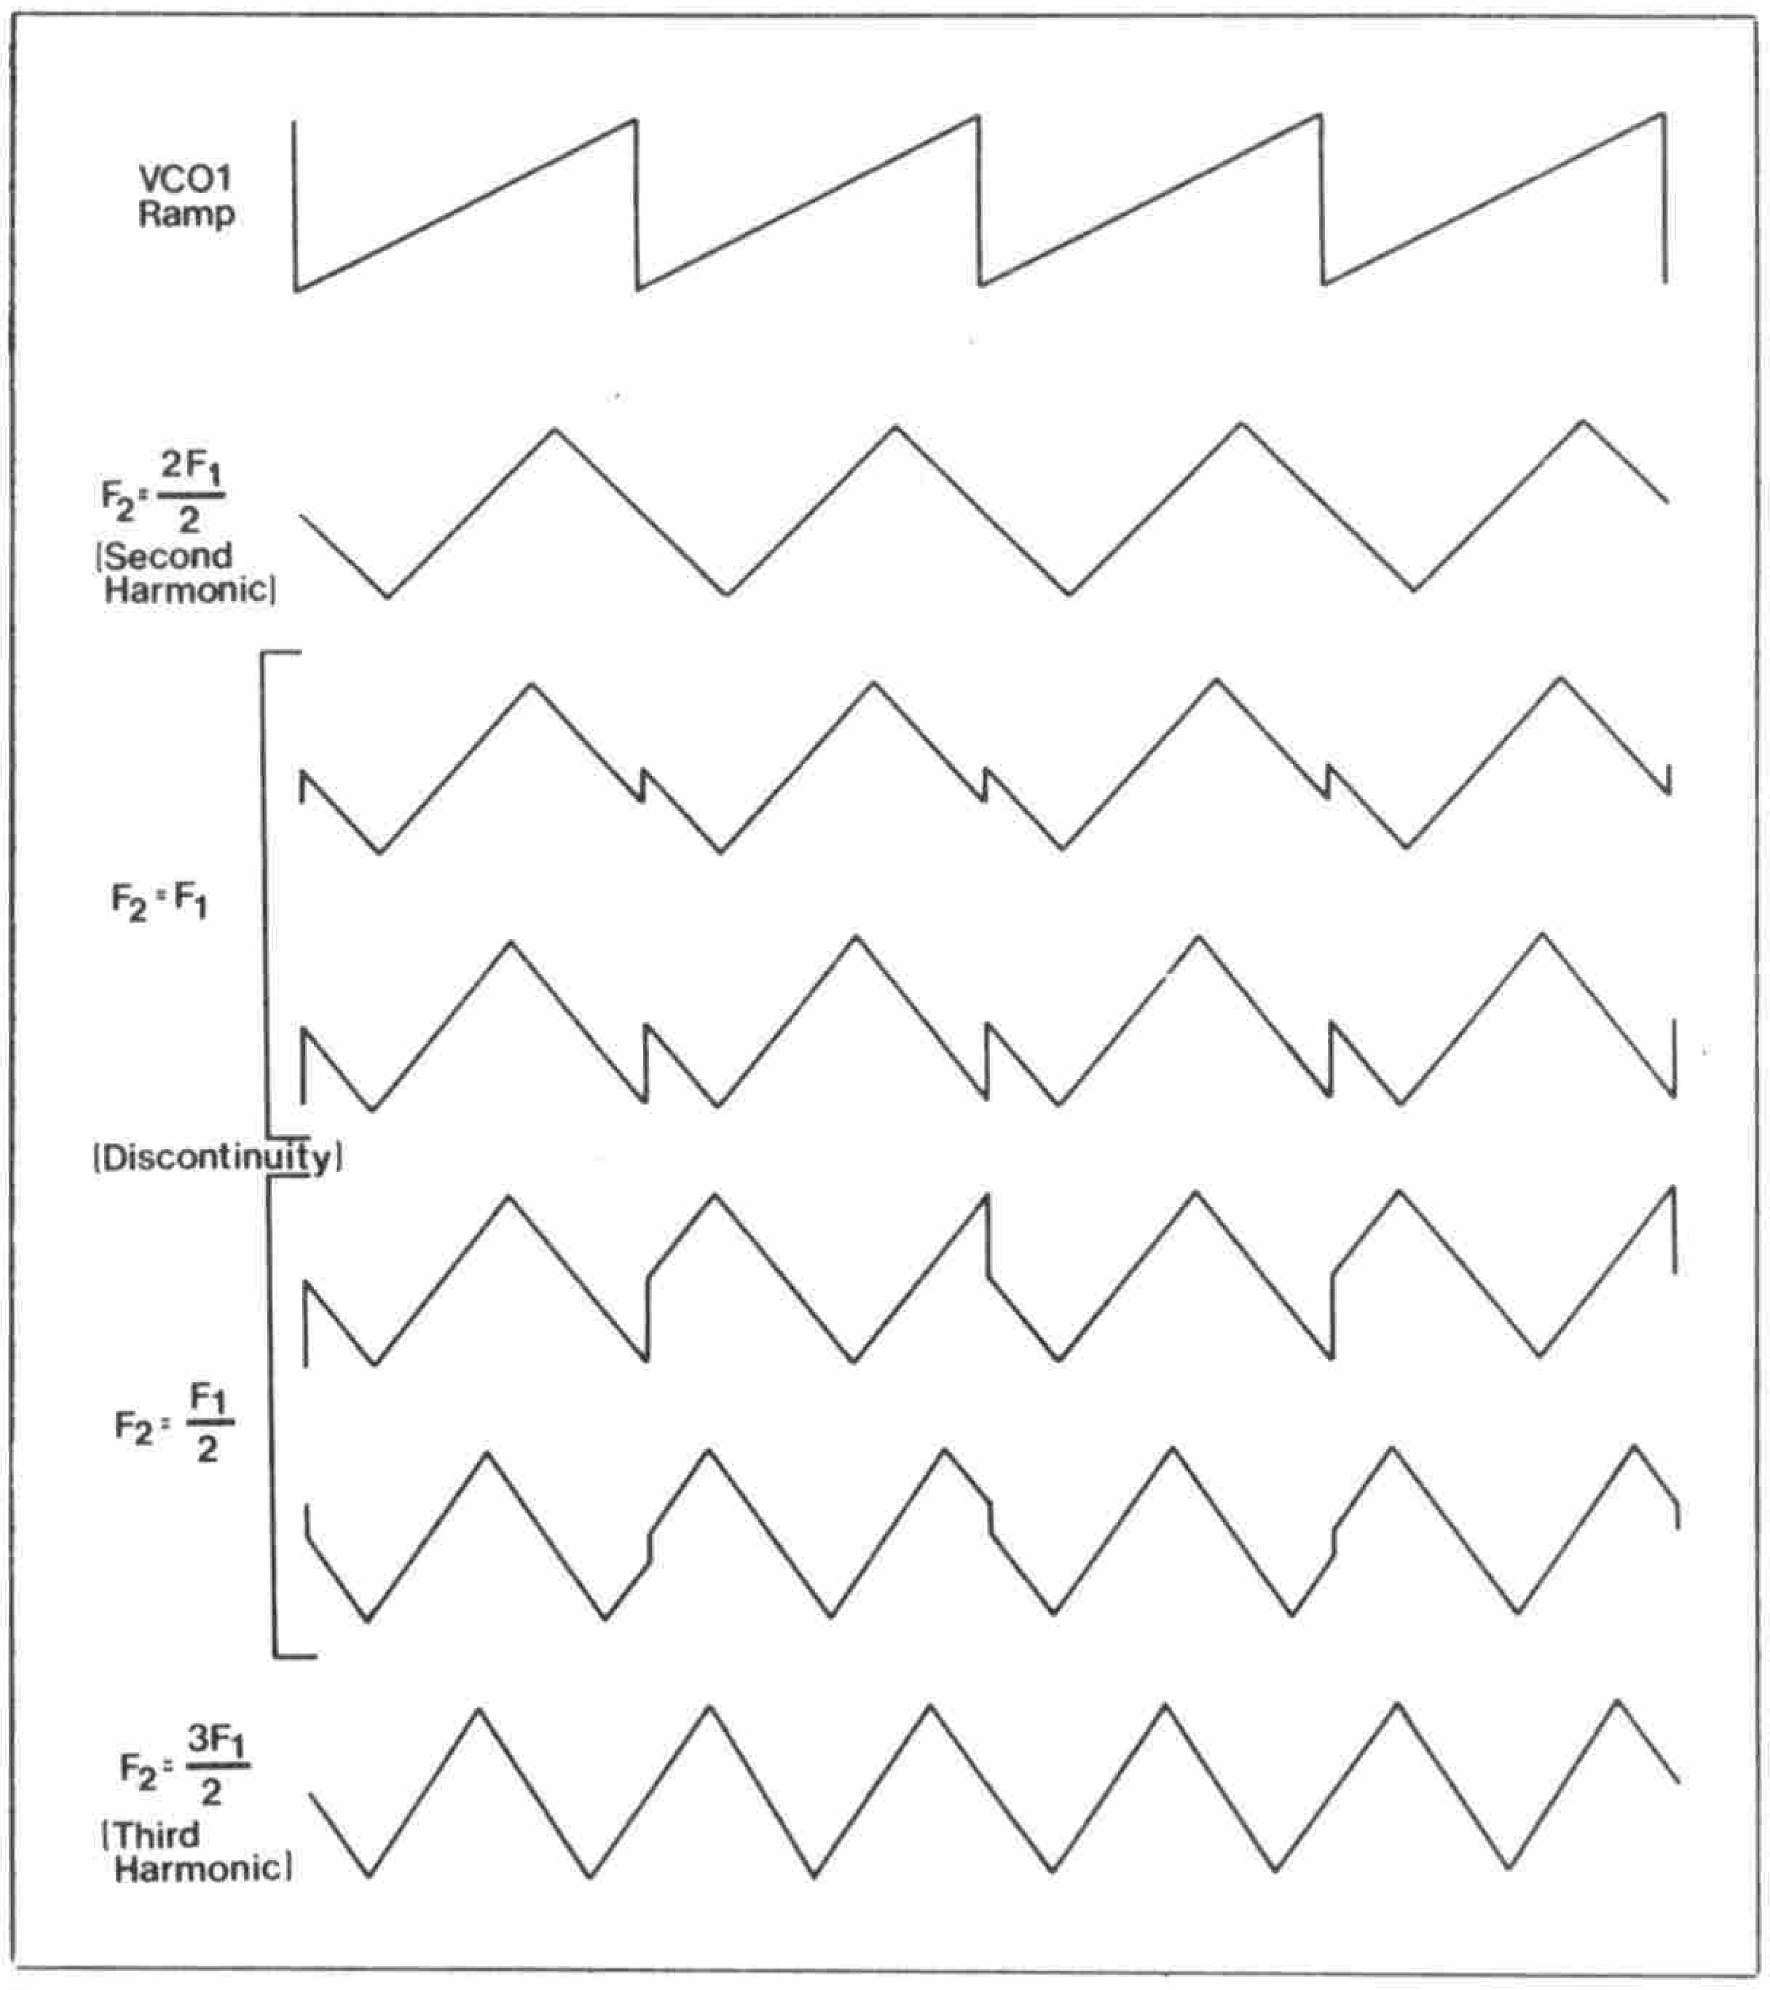

Figure 5. Sync II.

threshold is reached. The buffer produces an offset of about -1.6V and since the comparator refers to the output of the buffer, the voltage on the capacitor ramps between approx +1.6V and +6.6V. RV30 is set to return the buffered waveform to just below 0V, corresponding to about 1.6V on its wiper. This makes sure that the internal comparator is set to its rising state by the waveform crossing the lower threshold. Hence Sync I causes the triangle wave to begin an upward slope from its minimum value at the end of each VCO1 cycle. Sync II differs in that the triangle wave is set to its midpoint and proceeds in the same direction as before the sync pulse, i.e. the comparator state is unaffected. This means that slight changes of frequency that bring the VCO2 triangle wave to a peak before the sync pulse, where the sync pulse previously caught the waveform just before it reached the top, cause discontinuities in the tone and pitch of the sound. This is a feature of the pitch quantizing effect of Sync II, where the pitch of VCO2 jumps from one harmonic to the next as the control voltage to VCO2 is increased. As a result of the fact that alternate sections of the waveform between sync resets are inverted if the sync occurs on alternate rising and falling slopes, there is an inherent divide-by-two so the harmonics generated are really those of the sub-octave of VCO1. Figure 5 shows some examples of Sync II waveforms, those between the second and third harmonics of the sub-octave of VCO1. Note that as the rate of VCO2 is increased harmonics of the VCO1 fundamental increase in amplitude until the period is suddenly doubled with the introduction of the sub-octave component and from there on the harmonics diminish until the triangle wave is restored at a higher frequency.

Sync I produces a smooth change in timbre as VCO2 is swept, since each time sync reset occurs, the cycle starts in the same way. This makes it more useful for timbre modulation, whereas Sync II is best for pitch effects. One of the simplest uses is to generate the effect of a full-wave rectified ramp wave which can be modulated from a complete ramp to a triangle wave from the triangle output. On other synthesisers this is accompanied by a volume change, the triangle wave being half the amplitude of the rampwave, but with synchronisation the level remains fixed over the range, and of course the waveform shape can be swept much further in both directions. As the rate of VCO2 is decreased, a diminishing rampwave is produced giving a new method of amplitude modulation. As it is increased, the band of accentuated harmonics sweeps up the spectrum. Figure 6 shows some Sync I waveforms obtained from the triangle output with different relationships between the rates of the two VCO's.

So far we have only considered hard synchronisation, where the VCO2 cycle is restarted on every cycle of VCO. This gives the output of VCO2 the same period as that of VCO1, or in some cases of Sync II, double that. If the natural frequency of VCO2 is adjusted to a multiple of the VCO1 frequency, it will produce its natural waveform though beating effects are eliminated and a slight change of either frequency will introduce components of the VCO1 waveform into VCO2's output revealing the true period. Soft synchronisation causes reset only if VCO2 is past a particular point in its cycle and enables the pitches to be locked in musical intervals corresponding to fractional frequency ratios such as 3:2 (a perfect fifth), 4:3 (a perfect fourth) and 5:4 (a major third). Conventional discreet rampwave oscillators achieve soft sync by putting pulses on the ramp's upper threshold in the same way as the Spectrum LFO produces its regular S/H waveform. The Spectrum VCO's use a more advanced method which allows precise sync in ratios as low as 500:499 for example, where the VCO2 waveform is reset once every 500 cycles. Such weak synchronisation is heard as a series of clicks rather than an actual change of VCO2's pitch, but intermediate settings can give complex waveforms suitable for imitating many elusive sounds with complex harmonics such as those of engines, creaking doors etc. The synchronisation control in the FM & Sync section varies the depth of Sync I or II from zero (equivalent to no sync selected by the function switch) through increasing depths of soft sync to hard sync at the maximum setting.

The Synchronisation control uses the pulse wave facility of the CEM 3340 to inhibit reset until the rampwave of VCO2 has passed a certain point in its cycle. Reference to Figure 3 shows that the pulse wave is normally derived from the rampwave by comparing it with the voltage at pin 5, the pulse width modulation input. The output at pin 4 is an open NPN emitter, which is high while the ramp waveform is below the PW control voltage. This output is connected to the junction of R210,R211 in the base circuit of TR15 so for the first portion of VCO2's cycle the TR15 is held off and the sync pulses are prevented from resetting the cycle.

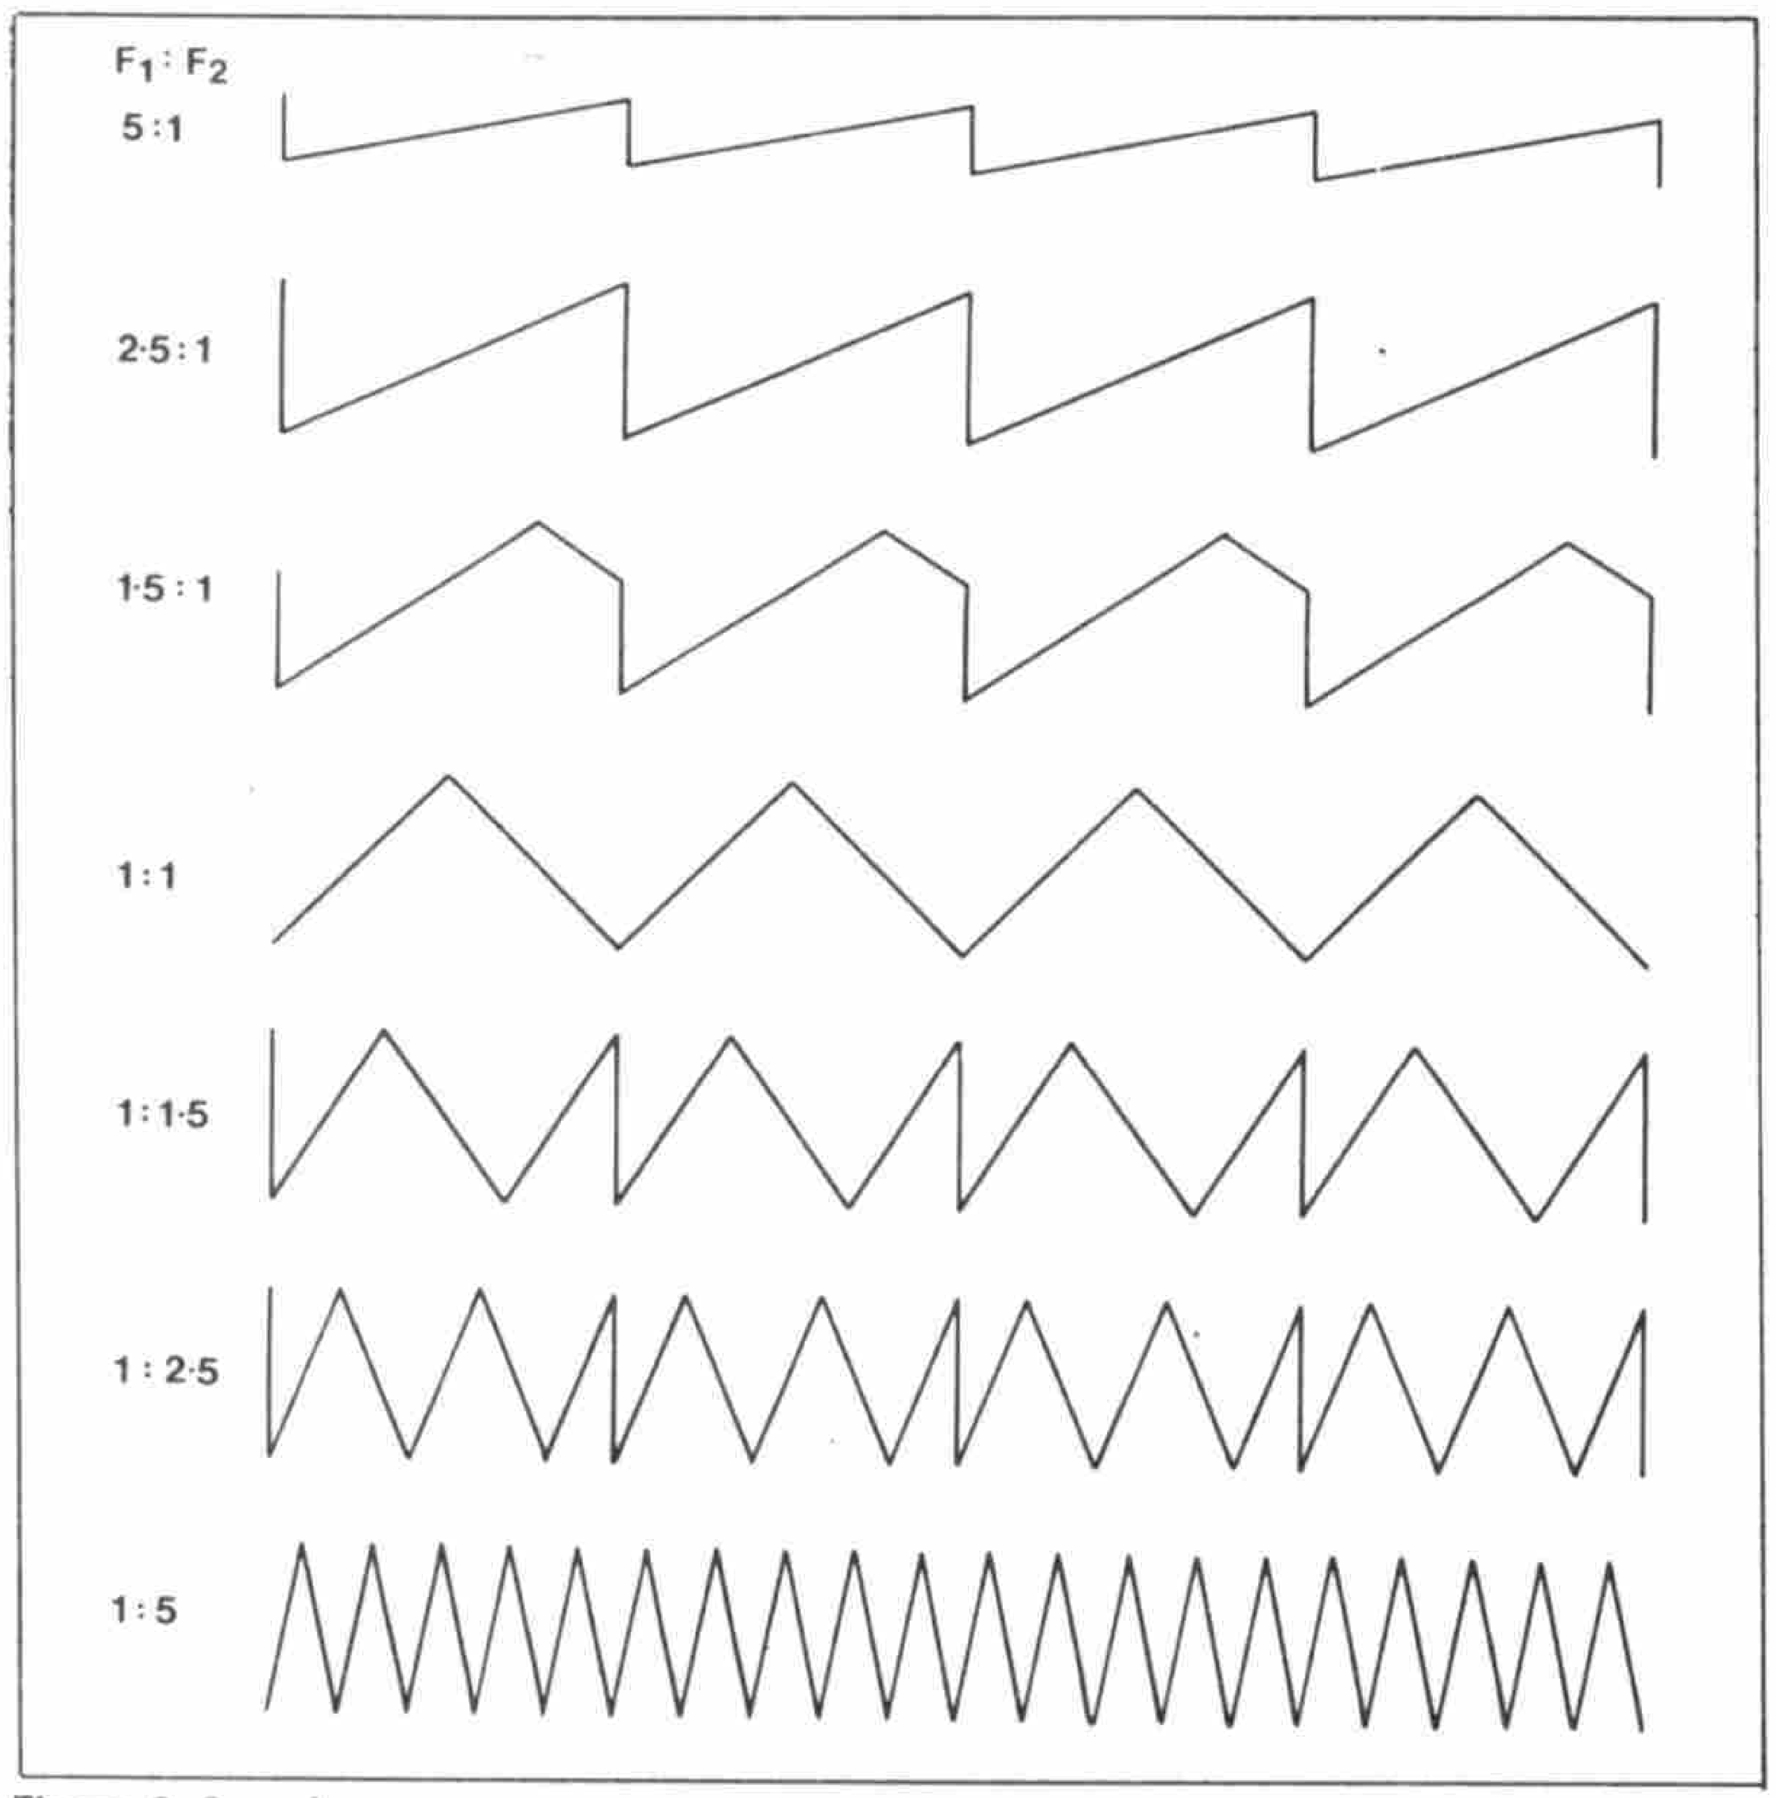

Figure 6. Sync I.

The proportion of the cycle for which sync reset is inhibited is determined by the setting of RV26, the Synchronisation control, which supplies a variable voltage to the PW control input. With the synchronisation control at 0 (>5V at pin 5) no sync reset can occur. At 10 (0V at pin 5) the PW output at pin 4 has no effect and every sync pulse causes reset (hard sync).

Figure 7 illustrates an example of how soft synchronisation (using Sync II) locks the pitch of the slave VCO (VCO2) in a musical interval with that of the master VCO (VCO1), in this case a fourth (a frequency ratio of 4:3). The sync pulses and waveform at the base of TR15 include positive going pulses (produced by the rising edges of the pulse wave) but these have no effect on circuit operation so are omitted for simplicity. Without the synchronisation operating, the ratio of the VCO frequencies would be 39:30, a flat 'perfect' fourth. The dotted line shown against the VCO2 ramp wave represents the level at the PW control input of IC15, pin 5, and corresponds to a setting of 3 on the Synchronization control. While the ramp is below this level the base of TR15 is held high, blocking the sync pulses. The phase of the higher frequency VCO2 waveform advances until a sync pulse coincides with the portion of the ramp above the dotted line, and the VCO2 waveform is reset to zero (point 'A'). This brings the ramps into phase, and until the sync pulse is again successful (point 'C'), VCO2 runs freely. Though the fourth above VCO1's pitch is heard clearly in the output of VCO2, the actual pitch of VCO2 is two octaves below this, at the lowest common denominator of the two frequencies. In practice this can be eliminated by tuning VCO up until a near perfect triangle or ramp is produced, with the sync pulses just catching the end of each fourth cycle, but since the extra components form a third note at the root of the chord it is often left in to produce richer sounds.

When using soft synchronisation, the PW output of IC16 turns TR15 off as soon as the reset takes the ramp waveform below the voltage on the wiper of the sync control (the dotted line). This would cause the new cycle to begin at some point above 0V (or with Sync II above 2.5V) depending on the point it was at before the sync pulse. C38 is included to keep the FET on for a short time after the reset turns TR15 off, ensuring that C34 discharges to the voltage on the potential divider.

The pulse wave output of VCO1 is variable from 0 to 50% by the Pulse Width control and from 0 to 100% with modulation. IC14b sums the voltages from the PW control, modulation routing and controller. The output is low-pass filtered by R180,C28 before being fed to the PW control input of IC15. This is to prevent stray feedback from the pulse output causing a fast burst of pulses on the falling edge which would confuse the sub-octave generator. C29 performs the same function on IC16, preventing spurious synchronisation pulses. The pulse output at pin 4 of IC15 is pulled down by R187 (being an open NPN emitter). The waveform is sharpened up by IC17a, a Schmitt NAND gate connected as an inverter, and used to clock the flip-flop IC18a. This produces a square wave of half the frequency, which is mixed with the pulse wave to give the sub-octave waveform.

The flip flop input would oscillate with the slow edges of the raw pulse output of IC15, so the schmitt gate is necessary for proper division.

The pulse output of IC15 is the source for the synchronisation circuit, so the sync effect can be turned on and off by modulating the PW through 0%. Hence, for example, the joystick can be used to bring in parallel harmonies or the free phase sound of unison oscillators could be introduced by the envelope generator as a note decays.

The VCO2 pulse is derived from the rampwave by IC17b. RV27 allows its width to be set between 0 and 65% at the setting up stage. VCO2's sub-octave waveform is generated by IC18b.

The ramp wave outputs of IC15 and IC16 are used directly and the triangle wave outputs are buffered by IC19a and b respectively. The half-way rectified ramp waveforms are produced by mixing the triangle and ramp waves in equal proportions. S9 and S10 are the waveform selector switches for the two oscillators, and connect to a virtual earth summing node in the VCF circuit. R195-208 are chosen to give equal peak amplitudes for the different waveforms.

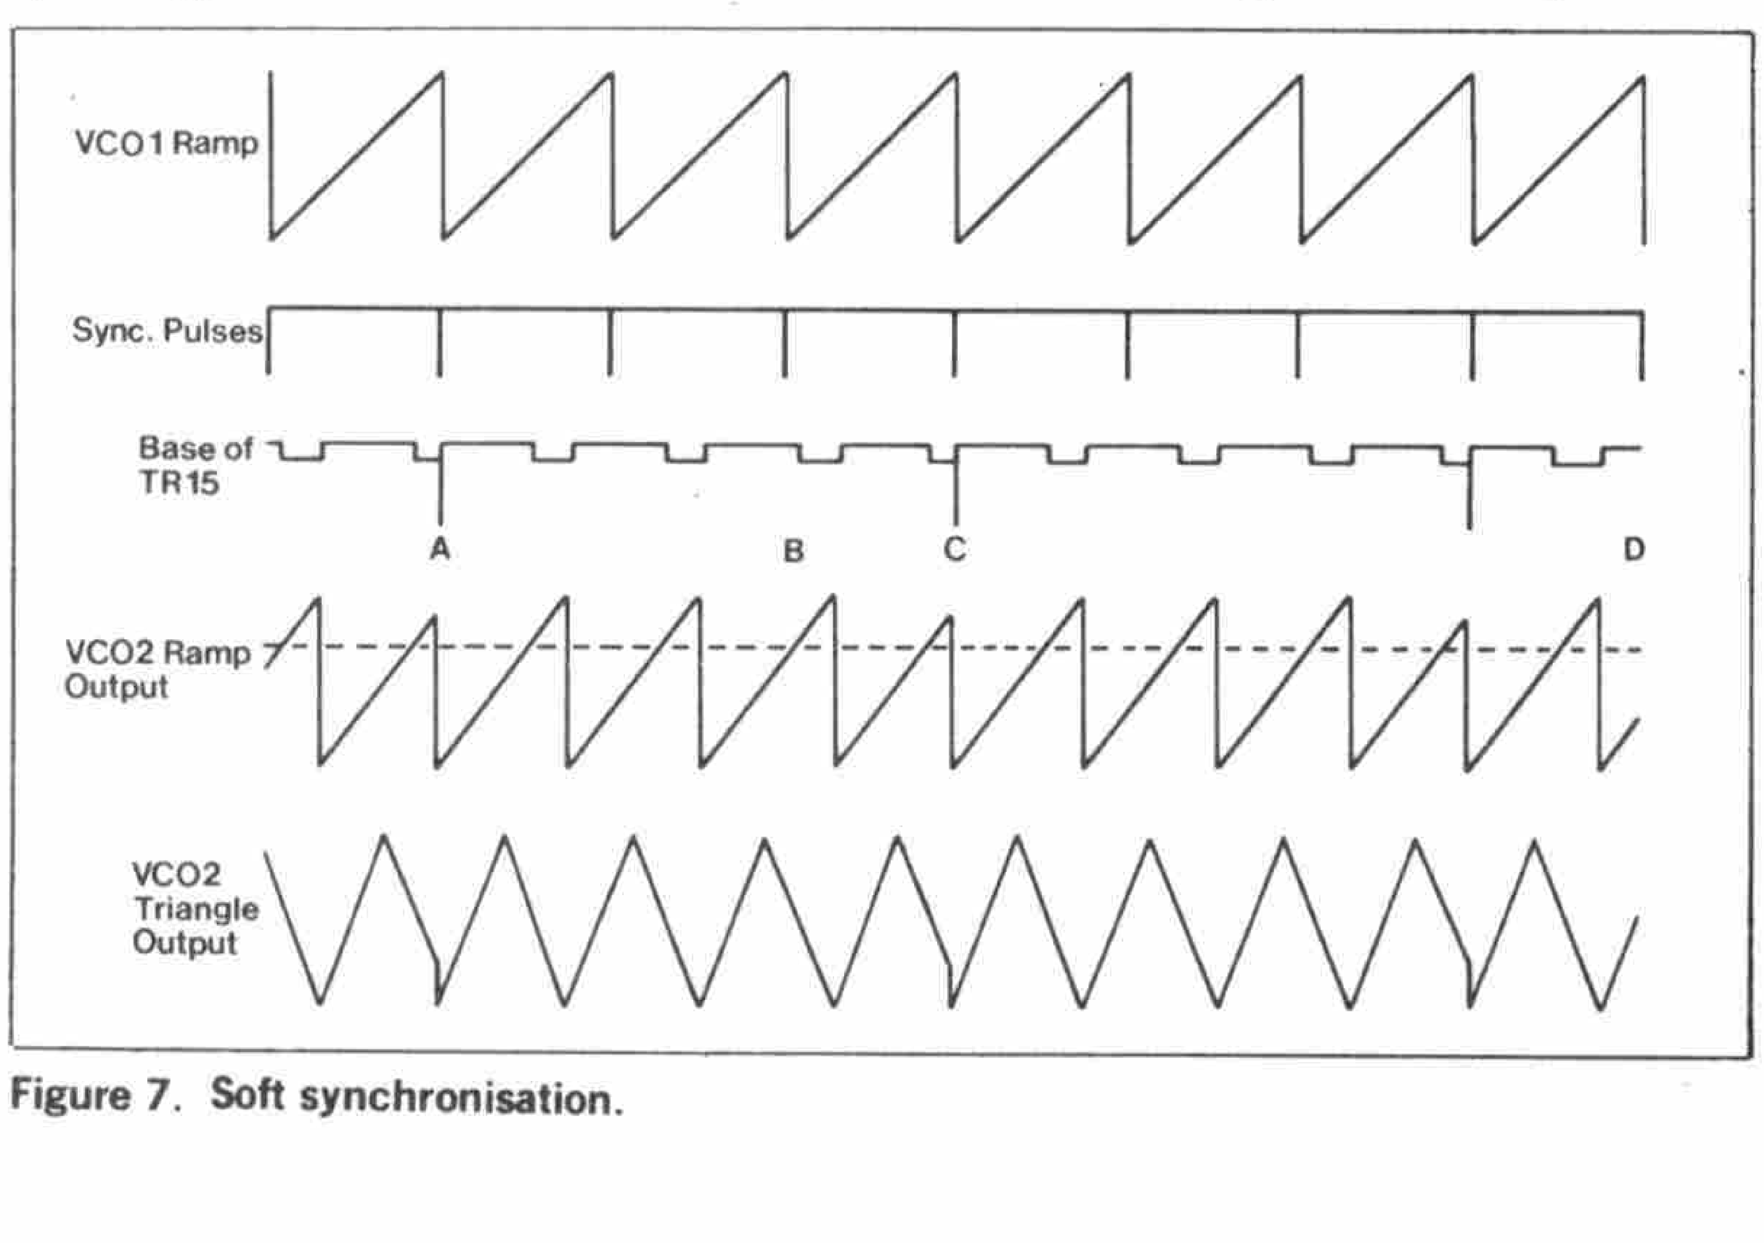

Figure 7. Soft synchronisation.

The two sub-octave square waves are NAND-ed to provide the drive to the tuning LED. When the waveforms are out-of-phase, the output is high and the LED off. Advancing phase difference due to slightly different frequencies produces a pulse wave that varies from 100 to 50% width, displaying the beats as fluctuating LED brightness.

The ring modulator is based around IC20 and processes the pulse wave of VCO1 and the triangle wave of VCO2 to produce complex non-harmonic sounds. It functions in a similar way to the ramp wave shaper of the Spectrum LFO by inverting the triangle wave about its midpoint when the pulse wave is high, and leaving it unchanged when low. This constitutes four quadrant multiplication of the value of the triangle wave by the value of the pulse wave (-1 or +1). When the pulse output is low Tr12 is off and the triangle wave is inverted with a gain of 2 by IC20a. The output is mixed with the original triangle wave of half the amplitude and opposite phase by IC20b. With the pulse output high the collector of TR12 is at -15V and the output of IC20a is positive. This reverse biases D32, and no signal reaches IC20b via R221. The original triangle wave is inverted by IC20b and shifted by the current through R220. The output of IC20b is the required product.

The noise generator is quite conventional, using the thermal noise of a semiconductor junction as a source. TR14 amplifies the noise on the emitter of TR13 to about 4mV p-p, which is boosted to ±2.5V by IC21. RV31 mixes the noise and RM signals and RV32 controls the amount sent to the VCF section. The noise signal is also sent to the LFO and modulation routing sections.

Part 4 will describe the remainder of the circuits, including the VCF, VCA's and Envelope Generators.

Since the Spectrum article began, we have had many enquiries from readers wishing to begin construction immediately. We strongly advise waiting until Part 5 is published, by which time the main parts list and full construction details will have appeared.

E&MM